

Running TUFLOW on the Cloud

Author: Jaap Van Der Velde and Chris Huxley
November 2021



1	Background	1-2
1.1	Context	1-2
1.2	TUFLOW Execution	1-2
2	Benefits of the Cloud	2-4
2.1	Varying Compute Demand	2-4
2.2	Physical Limitation to Available In-house Hardware	2-4
2.3	Unlimited Compute Resources	2-5
2.4	Access to Latest Technology	2-5
2.5	Location Independence	2-6
3	Challenges	3-7
3.1	Lifecycle Considerations	3-7
3.2	System Design Considerations	3-7
3.3	Orchestration	3-8
3.4	System Maintenance	3-9
3.5	Licencing	3-9
3.6	Access Control	3-9
3.7	Cloud Cost Management	3-10
4	Solutions	4-11
4.1	TUFLOW's Free Demo Mode	4-11
4.2	Free TUFLOW Benchmarking	4-11
4.3	Changes to TUFLOW	4-11
4.4	TUFLOW licensing	4-12
4.5	Steps to Run TUFLOW on the Cloud	4-13
4.5.1	Remote Desktop Connection	4-13
4.5.2	Network Connectivity with your Existing Network Licences	4-14
4.5.3	Setting up a Licence Server in the Cloud	4-15
4.6	Storage	4-16
4.7	Automation	4-16
4.8	Alternative Simulation Service Options	4-17
	Appendix – Wibu CodeMeter Cloud Licence Server Advice	4-19

1 Background

As the move to cloud computing is becoming more ubiquitous, there is an increasing sense of urgency for many organisations to look to the Cloud as an alternative to local hardware solutions. However, be it for small businesses or large enterprises, adopting cloud computing isn't a trivial task. How can you separate the hype from the drivers that matter to you? What challenges can you expect on the way? What challenges have already been solved and what solutions are available?

1.1 Context

In the conversation about cloud computing, the waters can get muddied by speaking about 'public vs. private cloud', 'hybrid cloud', etc. In this document, we're using the concept 'Cloud' to just mean public cloud infrastructure and services such as Microsoft Azure, Google Cloud and Amazon Web Services (AWS). For sufficiently large enterprises, private cloud solutions may offer similar benefits, as usage and costing of these facilities is similar to public cloud usage for smaller organisations, but the finer distinctions are beyond the scope of this paper and do not have a major impact on the central questions.

As the developers and users of TUFLOW, we're approaching simulation on the cloud specifically from the perspective of running TUFLOW Classic, TUFLOW HPC and TUFLOW FV within such an environment. Most of the discussion applies equally to other types of modelling software that use a similar mode of execution:

- Individual simulations run by a central executable on a single machine;
- Primarily file-based inputs and outputs (in the order of several gigabytes up to hundreds or gigabytes over the duration of a project);
- Longer running simulations (in the order of hours up to several days).

1.2 TUFLOW Execution

The TUFLOW End User Licence Agreement was updated in 2018 allowing companies to host their own licences on the Cloud. The only restrictions associated with users running TUFLOW simulations on their own company public or private cloud environment are:

- The licence must be a "Network" type.
- Usage of TUFLOW software on a Virtual Machine (VM) is confined to Authorised Users within the Licensee's Network. This clause means companies cannot on-sell access to TUFLOW licences hosted in the cloud or otherwise (excluding TUFLOW vendor contract arrangements).

Configuration of your cloud environment is your own responsibility. There are numerous ways TUFLOW licensing and simulation can be configured in a cloud environment depending on the cloud provider and internal company protocols. We recommend engaging a professional with suitable cloud architecture expertise to design your bespoke system. Clients who have already migrated to the Cloud have done so in a variety ways:

- Some use a hardware lock (USB) dongle that resides in their office on a physical computer or server. Cloud VMs obtain the network licence through a network link between the Cloud machine and the machine with the hardware lock.
- Others use a software lock. Software locks are a digital licence file that is locked to a dedicated host computer, server or VM. When using a software lock please select the host carefully as the software licence will be bound to it. Relocating the licence to a new location will require TUFLOW sales staff to reissue the licence, which incurs an administration fee. Note that a Cloud machine rebuilt from a disk image constitutes a 'new location' for the licence, so keep the configured VM around in 'stopped' state if not being actively used.

This report discusses some of the benefits, challenges and solutions those who are investigating transitioning to the cloud for TUFLOW modelling should be aware of. It is worthwhile noting that in addition to that simulation option, [TUFLOW's Cloud Simulation Service](#) and [Jacobs Flood Cloud](#) are both operational TUFLOW cloud simulation services. These services are however not the focus of this document.

2 Benefits of the Cloud

Although cloud computing providers would have you believe cloud computing is always the best solution, we found it offers significant benefits in specific cases, while other cases favour owning your own infrastructure. Some of the beneficial situations are summarised below.

2.1 Varying Compute Demand

One of the main benefits of cloud computing arises from its ability to scale on demand. This is beneficial for organisations who experience large variations or volatility in compute demand. Common situations known to cause this volatility are:

- Natural variability in project workload;
- Provision of emergency response services that needs to respond to flood events with forecast modelling;
- Unique modelling approaches which can translate to mass simulation demands interspersed by periods with little or no runs. This is common of ensemble and Monte Carlo simulations in the hydraulic modelling industry.

Cloud computing offers the benefit of quickly provisioning compute facilities and making these available for simulation, while only charging you for the actual usage, provided you pick the right plan with your cloud provider.

A consideration with this benefit worth noting is that on-demand cloud instances typically cost substantially more than reserved instances. For example, at the time of writing, both Microsoft Azure and AWS charge down to roughly a quarter of the cost per time unit for long term reserved instances, compared to the same hardware under an on-demand agreement (see <https://azure.microsoft.com/en-au/pricing/reserved-vm-instances/> and <https://aws.amazon.com/ec2/pricing/>). As a result, it still pays to consider your utilisation rate and whether you truly need to vary your computing need as strongly as you'd like.

2.2 Physical Limitation to Available In-house Hardware

Each organisation has a physical limit to the volume of computer hardware available within an office or on a company network. We have found consideration of this varies depending on the hardware required for the compute.

TUFLOW Classic and TUFLOW FV both rely on Central Processing Units (CPU) for their compute. Due to the natural evolution of computing over the past few decades, CPU hardware is typically plentiful, both within an office environment and on the Cloud. If/when in-house compute resources are insufficient there is a huge range of different CPU options available in the Cloud to meet specific project needs in terms of compute speed, RAM and storage. These get competitive with in-house hardware when considering total cost of ownership quite quickly.

TUFLOW HPC can run up to 100 times faster than its regular mesh CPU alternative, TUFLOW Classic. TUFLOW HPC however requires Nvidia Graphical Processing Units (GPU) to fully achieve this runtime benefit. For many users, GPU hardware may be a limiting constraint within the office environment. The Cloud is a useful resource to address the in-house GPU limitation. Suitable GPU hardware is however not as prevalent on the Cloud as CPU, and as such there are far fewer suitable options by comparison. This is discussed further in Section 2.4. It's worth noting that at the time of writing GPU hardware on Microsoft Azure and AWS had halved in price in the previous year in the Australian Region (where the authors are based) and new hardware has since become available that may offer an even better price/performance ratio.

Another factor to consider, favouring cloud compute, is that hardware for the business and consumer market often is not designed to be put under load 24/7, or even a substantial part of that time. Server hardware, on the other hand, is designed to take this type of pressure and cloud providers will ensure their orchestration services spread the load to get the most out of infrastructure, so you don't have to.

2.3 Unlimited Compute Resources

Another perceived benefit of the cloud is access to seemingly unlimited compute resources. Although it is true that the Cloud provides quick, on-demand access to amounts of computation you would typically never want or be able to purchase, they are far from unlimited. Depending on the specific cloud provider you select, they will provide access to on-demand resources within set quotas, for different types of resources. These quotas can be raised, but of course this is limited by how much the cloud provider is willing and able to commit to making available to any single subscription, at any given time. The TUFLOW team have experienced hitting the upper quota for on-demand GPU cloud hardware!

For some providers like AWS, this equation will be further complicated by the concept of spot pricing, where you get the resources when they happen to be available, for a maximum price you're willing to pay. But they will be taken away if someone who is willing to pay full price needs them! As this happens, your instance will be hibernated, stopped or terminated. It's worth noting that TUFLOW Classic, HPC and FV have been designed to deal with this type of interruption via the use of write/read restart files. For example, any tracked maximums from a simulation used to write a restart file will be retained and accounted for by the subsequent read restart simulation. If the use case is compatible, configuration of a model with periodic write restart files can help safeguard against unexpected simulation termination. Otherwise, it would only serve you if your simulations are relatively short, and you don't mind having the odd simulation prematurely terminated.

More generally, it is recommended to manage your quotas so they match your actual needs, which will make the cloud provider more amenable to accepting requests to raise them, provided they have the capacity to meet the need.

2.4 Access to Latest Technology

A definite benefit to the Cloud is access to the latest technology as it becomes available. By comparison, computer hardware you purchase will typically have to serve you over whatever period you depreciate

it, effectively locking you in to a minimum of 3 of 4 years. Thanks to healthy competition, cloud providers can provide you with access to the latest and greatest, as well as any hardware they still have available from before, for a more competitive rate.

There are however some caveats. The exact hardware that's optimal for your particular simulation, from a cost optimisation perspective, may not be cost effective or even practical for the Cloud provider. A good example of this is cloud GPU hardware. Cloud providers typically offered NVIDIA GPUs with Tesla (Pascal/Volta) architecture, which are double precision server work horses. Meanwhile GPUs with the Turing architecture (like the 20-series GeForce RTX gaming GPUs) which excel in single precision computation, are substantially cheaper, and more than sufficient for TUFLOW HPC simulations. As such, although the Tesla NVIDIA GPUs on the Cloud are suitable for TUFLOW modelling, such a high grade GPU would typically not be necessary. Cloud providers also tend only to offer GPU instances paired with a large numbers of CPU cores and massive amounts of memory and storage, often far exceeding what your TUFLOW modelling may require. The typical scenarios these Cloud GPU instances are designed for (machine learning, etc.) require these resources, but your simulations may not so effectively you end up paying for hardware you don't need.

In 2021 and beyond, the numbers are set to change somewhat as the 30-series RTX and server-based A100 series share the Ampere architecture, but at the time of writing it is too soon to say if this will invalidate the above. The majority of hardware available currently is still Tesla and Turing hardware.

2.5 Location Independence

Relative independence of location is arguably a main strength of the Cloud, allowing you to access resources from any location, without having to worry about where your data and infrastructure lives exactly – as long as you keep any formal restrictions regarding national borders in mind, quite often some clients will have restrictions on where their data can be stored and processed.

However, where you generate and consume your data is an important logistical consideration. If you do business internationally generating results closer to where they need to be consumed, or where they need to be stored for an extended period can be more economical than having to transport the data. For small volumes, this tends to be trivial, however for volumes that some types of project will consume or generate (sometimes exceeding TBs), you would do well to consider it before selecting a provider and a region. Also, depending on the nature of the simulation, latency to your services or external resources may also figure into your choice of region.

A clear benefit of cloud infrastructure with most providers is that they offer you otherwise identical services in several regions across the world, close to wherever you need to deliver or keep the data.

3 Challenges

3.1 Lifecycle Considerations

The freedom to select Cloud resources on a platform and in a region that most benefits your use case is one of the strong points for using cloud compute – but making that choice can also be a major challenge. All major providers do offer a wide selection of both Windows and Linux images and instances, as well as any supporting services you are likely to need. Not all providers are equally well-represented in all regions. When considering setting up a cloud environment with a specific provider and in a specific region, consider your process from start to finish, the entire lifecycle of what you're looking to produce. Simulations that periodically run unattended and continually deliver data to critical processes (such as flood forecasting systems) require an entirely different environment from simulations that run in a more project-like setting with large numbers of simulations, re-runs and on-site analysis, resulting in a research outcome or commercial study. Some useful questions include:

1. What are your simulation software's requirements?
2. What data are your simulations relying on? Do you need to keep them on hand, or are they small enough or changing so quickly that you need to provide them with every started simulation?
3. Where are results collected and where do they need to be delivered? Have you considered pre- and post-processing of data and can or should this be done in the Cloud as well?
4. What level of access do you require during simulations and in between simulations, when you're running a larger number of them? How will you monitor progress and how do you determine simulation success (or conversely how do you easily identify simulations that have failed)?
5. Are your simulations or processing tools dependent on services or resources elsewhere in your infrastructure, like service APIs, licence servers, etc.?

3.2 System Design Considerations

TUFLOW Classic and HPC is currently only available on Windows operating systems. It processes its inputs, outputs and logging purely as files, so it requires no access to external resources and offers no monitoring other than through those outputs and logging. TUFLOW FV is similar though is currently available on Windows and Linux operating systems.

Although inputs for TUFLOW simulations can be large, simulation outputs tend to be far larger, by orders of magnitude. Many simulations share inputs and the software optimises consecutive runs by caching previously parsed inputs in binary form on disk. Because of this, and since simulation runs are monitored through their file outputs, allowing the simulations access to a shared location for outputs has management benefits. As such, in the majority of cases TUFLOW runs entirely in the cloud utilising

cloud storage. In specific use cases, shorter running independent simulations or those depending on variable inputs like forecasting simulations are run as independent tasks with separate up-/downloads.

Larger simulation projects often warrant setting up separate remote desktop VMs in the cloud that can run alongside the storage. The VM can be used for input preparation tasks or result post-processing. The added overhead of setting up and configuring a virtual desktop machine with the appropriate tools is often more time efficient and less cost and effort than up-/downloading results for analysis on a standard desktop environment. Further to this, working on the Cloud directly can help reduce the size of the data that needs to be delivered or stored long term. In flood forecasting cases where the inputs for a project are obtained from other online sources, it may even be worth considering setting up an entire project environment in the Cloud, with all data staying within the environment for the project duration.

Network traversal is a limiting factor for larger projects in several ways. Directly accessing files on the cloud from a desktop in your organisation's local network can be relatively slow. Most cloud providers charge for any data leaving the environment, even if that's only a read operation to 'look at' a file; however, the cost of downloads and remote reads is generally insignificant when compared to the overall cost of compute and storage. There may be more serious costs associated with excessively large downloads that are easy to overlook, like exceeding a bandwidth cap or a monthly download cap from your organisation's internet service provider. Additionally, processing very large downloads can tax your organisation's network, causing delays for other networked activities. Finally, given all combined bandwidth restrictions, the total time needed to download results may affect lead time substantially. Major cloud service providers offer alternatives for massive data egress, including simply shipping the data on hard drive from the data centre.

Whether download or upload poses the greater challenge depends on the type of modelling and your specific model configuration. It's also worth noting that moving data into a cloud environment is typically free (you will pay for storage once it's there), while getting the data out, or moving it into another region, has a charge associated with it.

Another data-related cost and usage caveat: storage in the cloud is associated with safety and automated backups. However, not all cloud-based storage services have backups included in their price and in many cases you have to explicitly configure backup services through management interfaces – and pay for their usage.

3.3 Orchestration

Once fully configured and running, a cloud environment may be very similar to a more traditional computing environment. Cloud providers offer specialised tools for provisioning and managing resources that you should consider when taking your modelling and simulation workflow to the cloud. A simple 'lift and shift' approach may be a good starting point to acquaint yourself with the cloud, but facilities like batch processing and automated provisioning of resources can greatly increase the benefit which the Cloud provides.

Different cloud providers provide very different sets of tools, with their own Application Programming Interfaces (APIs) and Graphical User Interfaces (GUIs) to control them, for provisioning, orchestration and management of your cloud environment. As a result, relying on those tools carries the risk of “lock-in” to that provider. If you need to be able to run simulations across different cloud platforms, keep in mind that the environment your simulation is running in can be created to be fairly uniform across platforms, with operations specific to each platform clearly isolated from the generic workflow.

Due to the complexity and extent of cloud infrastructure tools and environments, ultimately providing environments that are surprisingly simple and familiar to use, it seems likely that ‘simulation as a service’ will see a rise in popularity. An important consideration is whether the complexity of your project/s removes the need for the skills required, makes it worth investing in them, or moves you to consider seeking a service provider.

3.4 System Maintenance

The Cloud is a new and evolving field. As such, we have found, the systems provided by the main cloud services providers, such as Microsoft Azure, Google Cloud and Amazon Web Service are still evolving too. As a result of this ongoing maintenance of any cloud system is required to maintain compatibility with the cloud service providers environment.

We recommend companies who are considering moving to the Cloud engage professionals with suitable cloud architecture expertise to assist in the application design. Further to that we also recommend the same professional be retained in a maintenance capacity, so they are available to fix your company system in the event changes made by the services provider compromise any automated workflows.

3.5 Licencing

Both the simulation software and other tools you need in the cloud may have licensing restrictions or require access to licensing services that restrict what you can do in the Cloud or how much effort and added cost will be involved. If you need to set up a licence server for a particular product, you’ll want to automate the process for quick repeated deployment or opt to run a licence server in the Cloud for the long term. Again, you’ll have to consider what region this service will be available in or arrange for good connectivity to other regions.

Although it’s true that cloud infrastructure is in many ways identical to physical infrastructure from a software point of view, be aware that some product licences can have special provisions on running in data centres.

TUFLOW specific considerations for licencing on the Cloud are provided in Section 4.4.

3.6 Access Control

Another non-functional aspect that can cause complications is access control. Many organisations haven’t advanced their use of cloud infrastructure to the point of connecting it to their owned managed infrastructure. As a result, directory services and all the policies that come with a well-managed network

and its users won't be available on the Cloud infrastructure. All major cloud providers provide excellent facilities to manage these aspects and, in most cases, you could integrate them with your existing systems, but this is not trivial.

If your modelling work benefits others, also consider their need for access to the environment as well as their expectations of privacy and security throughout the process. If the client's requirements exceed your own, consider running in an isolated environment, or write your tools and procedures so that they can be applied to the client's own tenancy in the cloud. This can also greatly simplify costing and invoicing, as you no longer need to forward the cost of spent infrastructure resources.

For TUFLOW simulations specifically, the running environment is typically separated from the inputs and outputs, making the entire solution relatively portable and access control a bit more straightforward than other situations.

3.7 Cloud Cost Management

Costing can be a challenge in either extreme. Without appropriate monitoring, entirely automated solutions can run rampant if problems go unnoticed for too long. The same flexibility and scalability that makes the Cloud so attractive can become a liability when design and execution isn't sufficiently controlled. In some situations larger projects may typically suffer delays or overruns that quickly increase cost in a more noticeable way than the same issues would in traditional infrastructure.

Additionally, when designing a solution that incorporates cloud infrastructure, it is as important as always to consider the total cost of ownership – cradle to grave. This is true when considering the implementation of your workflow in the Cloud, considering data ingress and egress and the various processing steps. But it equally applies when considering the long-term shelf life of your solution. Consider that both the underlying technology and the ways in which you access them through APIs and GUIs quickly develop and that obtaining infrastructure as a service makes you subject to the conditions and development of that service.

4 Solutions

Simulation in the Cloud is still a developing field. Some useful information focused on “solutions” is outlined below.

4.1 TUFLOW’s Free Demo Mode

A licence-free demo mode was introduced to TUFLOW in 2018. The demo mode is included within the standard TUFLOW executable, is fully enabled without the need for a paid licence, though has the following limits:

- 100,000 total cells and 30,000 active (potentially flooded) cells;
- 100 1D channels;
- There can only be one 2D domain (multiple 2D domains are not supported); and
- A simulation time of 10 minutes.

Clients wishing to migrate to the Cloud have done so successfully by first using the demo mode for their initial system testing and design. This meant they were able to optimise their cloud solution to meet their own specific requirements, before focusing on TUFLOW licence related challenges, also without using licences which their hydraulic modelling staff required for concurrent project work.

Example models using the demo mode can be downloaded from the Wiki: https://wiki.tuflow.com/index.php?title=Example_Models_Home_page

4.2 Free TUFLOW Benchmarking

Available hardware configurations from your selected cloud service provider may cause you to wonder which one would deliver the best performance. To help with this we have developed standardised benchmark models. The TUFLOW benchmark models are a licence free dataset that can be downloaded and used to compare different hardware options. The model does not require a licence to run a simulation. As such, testing using the benchmark model will have zero impact on any project work your organisation is undertaking using your paid-for licences.

The benchmark models and also results for a wide range of hardware are documented and available for download from: https://wiki.tuflow.com/index.php?title=Example_Models_Home_page

4.3 Changes to TUFLOW

Several updates have been made to TUFLOW in recent years that greatly improve the ease with which it can be used in a cloud setting.

- The TUFLOW Summary File (.tsf file) was implemented in the late 2000’s to allow users to obtain metrics from a model simulation while it is running. This is currently the recommended

file to read to monitoring simulation progress mid-simulation. It contains information on the solution scheme, hardware, primary simulation inputs, the simulation status, mass balance outputs, etc. Additional information will be added based on user feedback (please email support@tufLOW.com if you have a suggestion). The command *TSF Update Interval* can be used to control the interval in seconds to update the .tsf file while a simulation is running.

- Software licencing was introduced in 2016. Software licences are tied to a machine and do not require a physical hardware dongle. This licencing option was added primarily as an option so clients could host network licences on a Cloud VM.
- 2017 introduced a “no console” execution option. It disabled the console window when TUFLOW runs. As such, TUFLOW simulations will run, but no window will be visible to check model progress. This feature was built in primarily for running bulk simulations on blade machines, the Cloud, or on Linux platforms using a Windows emulator.
- 2018 introduced a feature to TUFLOW allowing users to obtain listings of all the files required during a specific simulation. This assists in the provisioning of resources and access to only include those specific files and locations, or to allow for smaller uploads prior to running smaller jobs.
- Licence free package model functionality was added in 2018. It provides an automated means to extract a complete copy of all input files for a given project, accounting for all possible scenarios and events, using a single batch file command. Many clients use this feature to package up model dataset for upload to the Cloud. Refer to the [2018 Release Notes](#) for further information.
- As of 2020, TUFLOW can run ‘headless’, which is to say it can be configured to be guaranteed to terminate, with a clear indication of success or failure – without any reliance on user interaction.
- TUFLOW Classic and HPC will be compatible with Linux in the future. This has a clear cost saving advantage on operating system licences on cloud instances, as well as opening the possibility to integrate with many tools available on that platform.
- Future releases also aim to include live logging facilities. Currently, TUFLOW allows the modeller to frequently write periodic status updates to a TSF file. However, since monitoring these files incurs a delay and may be challenging in a very distributed environment, TUFLOW will likely gain the ability to report a subset of logged information to network services directly.

4.4 TUFLOW licensing

TUFLOW relies on Wibu CodeMeter for access to a valid network licence. A physical licence (on a dongle) or a software licence is associated with a host computer, VM or server. Other computers or in the case of cloud computers, virtual machines are configured to source licences from the network host.

This process is relatively straightforward, there are however limitations with both options:

- A physical licence will be located within your own office in a desktop computer or server. Cloud VMs will require connectivity with the licence host to run a simulation. This will require a safe and appropriate network connection like a VPN or secured tunnelling, fitting with your existing IT.
- A software licence hosted on a VM in the Cloud will be bound to the specific instance it is first installed on. Once you install and activate CodeMeter and the TUFLOW licence, you won't be able to migrate it to another machine. As such, software licences are typically hosted on low cost CPU VMs which remain on indefinitely. Further to that software licences host on VMs are typically bound to the region you first activated it in. Use by VMs is permitted under the TUFLOW licence agreement, though will require configuring network access across regions, which is more complicated than hosting and running within the same region.

More details are provided in Section 4.5.

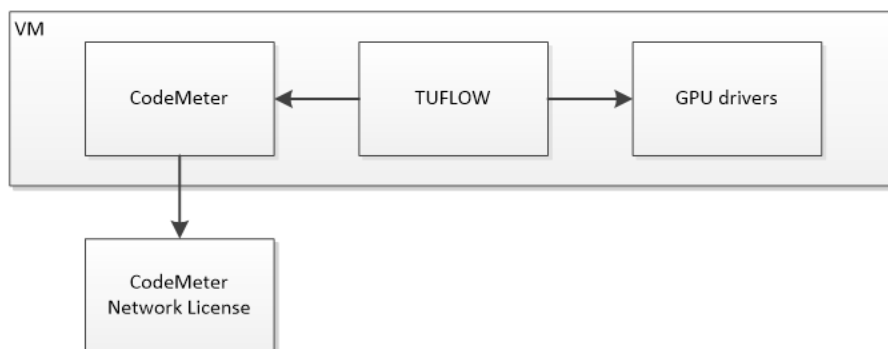
4.5 Steps to Run TUFLOW on the Cloud

Configuring TUFLOW in the cloud is relatively straightforward depending on how you'd like to run your models. Due to the wide array of different system configuration options available to meet specific client use cases the TUFLOW team do not actively get involved in setting up Cloud systems for our clients. We recommend you engage a suitably skilled Cloud professional for this task. Notwithstanding, the following sections provide some general guidance outlining a variety of available options.

4.5.1 Remote Desktop Connection

If you are just creating a VM based on a Windows image (any modern version of Windows, typically Windows 10 or Windows Server 2012 R2) and are logging on to it through remote desktop, installing and running TUFLOW is the same as on a regular PC


- Install and configure appropriate GPU drivers (for TUFLOW HPC), if they don't come pre-installed on the image;
- Install CodeMeter: <https://www.wibu.com/support/user/user-software.html>
- Configure Codemeter to access the network licence:
https://wiki.tuflow.com/index.php?title=WIBU_Configure_Network_Client_2016
- Copy your preferred version of TUFLOW to the VM:
<https://www.tuflow.com/Downloads.aspx>
- Configure and run your model on the VM.




4.5.2 Network Connectivity with your Existing Network Licences

When not using remote desktop, connecting a simulation on a Cloud VM to your network licence server in your office can be done in several ways. In all situations the configuration of Codemeter on the Cloud VM is no different from the standard configuration. Follow the steps outlined on the TUFLOW Wiki: https://wiki.tuflow.com/index.php?title=WIBU_Configure_Network_Client_2016. The connection to the network licence server in your office can however be configured in a variety of ways:


1. You can install a VPN client on the Cloud VM to connect to your organisation's network (if it provides a VPN service).



2. Alternatively, you can set up a VPN gateway with your cloud provider, but this will require setting up a virtual network, connecting the VM to that network and setting up the VPN gateway on that network. You can then connect to the licence server in your office using the IP it obtains when connected to the network in the cloud.



3. A third option is to expose the port on which the licence server (host) is running on your office network to be accessed from the cloud's IP. The VM can access the licence server on the company's IP and selected port number. Given the potential for abuse and relative lack of security, we don't recommend the latter option, but it will work if required.



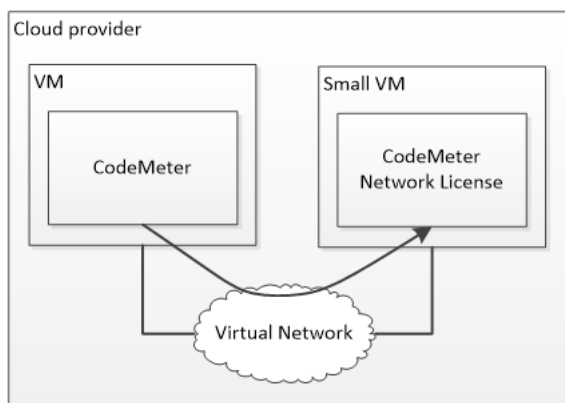
4.5.3 Setting up a Licence Server in the Cloud

For this configuration a dedicated VM will be used as the software licence server. Another VM will be used for the compute, though linked to the licence server VM. This configuration is recommended due to

1. The requirements for your simulation VM(s) most likely changing from one project to another
2. Compute VMs will also be more expensive to run than the licence server VM and as such will be an unnecessary expense if left open and unused while simulations are not running.
3. Software licences are locked to single VM. Physically relocating them incurs an administration fee from WIBU (please refer to the advice provided in the Appendix for more details).

The general system configuration steps are as follows:

- Set up a low-cost VM to act as the licence server (host).
- Set up a VM to act as the resource for running your model simulations (client).
- Create a virtual network with your cloud provider and configure both the licence server host VM and simulation client VMs to connect to it.



- Download and install CodeMeter on the client VM: <https://www.wibu.com/support/user/user-software.html>
- Request a software licence from sales@tuflow.com: https://wiki.tuflow.com/index.php?title=WIBU_Software_Licence_Update_Request
- Import the software licence to the host VM: https://wiki.tuflow.com/index.php?title=WIBU_Licence_Update_Import
- Configure the host VM as a licence server: https://wiki.tuflow.com/index.php?title=WIBU_Configure_Network_Server_2016
- Configure simulation client VMs to access the network licences on the host VM: https://wiki.tuflow.com/index.php?title=WIBU_Configure_Network_Client_2016

Note: Once the licence server is operational it is locked to the VM. If you remove and re-create the server, you will need to obtain a new copy of the licence which incurs an administrative fee. With this in mind, keep the configured VM around in 'stopped' state if you take it offline. Please refer to the advice provided in the Appendix for more details.

Typically, your virtual network will be local to a specific region for your cloud provider; creating connectivity between different regions may require a skilled IT network administrator. You can set up servers for each individual region, but you won't be able to share licences between the servers.

4.6 Storage

You may want to consider setting up storage in the cloud separate from your VM so you can mount it as a network drive in your VM. This way, if you write your results to it you can stop the VM once you're done running models but keep the data around in the cloud for an extended period without being charged for the usage of the VM.

4.7 Automation

Both to manage the complexity of a complicated and quickly changing environment, and to allow for reliable and repeatable results, automating the process of preparation, quality assurance, resource provisioning, orchestration, and logistics has major benefits.

Major cloud providers provide automation interfaces to shell scripting to match the operating systems you're running on, as well as exposing API's that are supported in several languages, of which Python is most common across providers. These API's allow you to remotely and securely operate the cloud environment from the inside and outside.

For example, if you are creating instances on an Microsoft Azure Batch account, you'll want to bundle graphics drivers, CodeMeter, and possibly TUFLOW in a batch pool application package (<https://docs.microsoft.com/en-us/azure/batch/batch-application-packages>). When creating a node, the application package will be installed. You can also consider putting TUFLOW in a task application package instead, which would be installed whenever a new task is started on the node, allowing you to run different versions of TUFLOW with new tasks.

4.8 Alternative Simulation Service Options

TUFLOW aims to continue developing its software to fit usage in this setting for all users, as well as providing running simulations as a service. If you are developing your own cloud system and have questions or alternatively would like particular features built into TUFLOW to assist please email support@tuflow.com.

If developing your own system is not achievable, two simulation services are available to assist.

1. TUFLOW Cloud Simulation Service (BMT)

This service is designed for bulk simulation of any TUFLOW product (Classic, HPC or FV). Software licences are scaled to meet your project demands (i.e. this is independent of any in-house licences you may own).

The BMT service offers flexibility regarding the Microsoft Azure Cloud hardware available for your project. We tailor the hardware and storage resource specification to suit your project requirements (based on project size, model size and the simulation type). The service has no limitations in terms of hardware type, offering simulation on a range of CPU and GPU virtual machines.

This service is well suited to mass production runs of an established model (often required for design event modelling or a Monte Carlo assessment). It is less suited to model build activities that require regular modeller/model interaction and iterations. We recommend model build activities be completed using a standard client licence as it is typically the most cost effective and time efficient option. More information on this service is documented on the [TUFLOW Wiki](#). Alternatively, email sales@tuflow.com.

2. Jacobs Flood Cloud

Powered by Parallel Works, Jacobs have developed Flood Cloud. Flood Cloud gives the modeller full simulation control on a cloud environment.

Flood Cloud includes a dashboard to assist users with all model uploads, simulation runs, simulation monitoring and result download tasks (requiring no external support from Jacobs staff).

Software licences are scaled to meet your project demand (i.e. this is independent of any in-house licences you may own). The service links to pre-defined CPU cloud hardware.

The Jacobs Flood Cloud offering is presently limited to TUFLOW Classic/HPC CPU based simulations. TUFLOW HPC GPU hardware and TUFLOW FV simulations are not yet supported.

For more information on this service please visit the [Jacobs website](#).

Appendix – Wibu CodeMeter Cloud Licence Server Advice

If you're setting up a Wibu CodeMeter licence server in the cloud, for use by TUFLOW, please keep the following in mind:

- As mentioned in Section 1.2, licence files are locked to a dedicated host computer, server or VM. A licence provided for a specific licence request will only work on the computer that request was issued from.

For cloud systems you will need to request (and pay for issuing) a new licence file if:

- You migrate the licence server to another region;
- You remove the licence server and/or rebuild it from its disk image (not all information that Wibu considers relevant to decide if a computer is “the same computer” is stored on disk).

Possibly also if:

- You upgrade the operating system to a new major version (next Windows Server, etc.);
- You substantially change the configuration of the VM, causing the cloud platform to effectively rebuild it for you.

We cannot be very specific here, in part because the licence server is not developed by TUFLOW, but by an independent third party. For the same reason, we cannot change the parameters of when CodeMeter decides your computer is now a ‘different’ computer.

- Although you can run the licence server on a server alongside other software and services, we recommend you don't. Instead run CodeMeter on a server that has just enough resources to run the CodeMeter service, according to (the fairly low) system requirements Wibu states for the version you install. Wibu CodeMeter has proven to be robust and is unlikely to cause issues by itself.

We recommend you maintain this dedicated server by:

- Periodically checking disk use on the server, in case of excessive logging, etc.
- Disabling fully automated updates if 24/7 availability is required, or schedule maintenance windows to keep your operating system sufficiently up to date and secure, depending on the update schedule recommended or required by your OS supplier / support organisation.
- VMs / compute nodes that need access to licences, need network access to ports 22350 – 22353 of the server. If you need to be able to deploy compute nodes in several regions, you either need to make the licence server accessible from the virtual networks those nodes are on, or deploy additional licence servers in those regions.

Network access to a licence server on the appropriate ports is all that is required to make use of the installed licences. Keep this in mind when securing your licence server and its virtual network.

- Setting up a CodeMeter software licence server does not remove the requirement to have a CodeMeter client installed on every compute node or VM you run TUFLOW on. The same installer can be used for the client and server role, but unless you run models on the actual instance that has the license installed, you need a client installed on the compute node and a server with an available licence. And the client has to be configured to know the hostname or address of the licence server, or the network has to be configured to make licence server auto-discovery possible.

Refer to Wibu's documentation for details: <https://www.wibu.com/support/manuals-guides.html>