CGMS best practices in support to using cloud for storage, storing, processing and dissemination of meteorological data

Endorsed by CGMS-52 plenary on 6 June 2024

TABLE OF CONTENTS

1. INTRODUCTION

This document presents the Coordination Group for Meteorological Satellites (CGMS) agencies Best Practices in support to using cloud for storage, storing, processing and dissemination of meteorological data. This document is developed by Cloud Expert Group, a sub-group under CGMS Working Group IV, established in July 2020 with the main goals of studying available technology, use cases of cloud utilisation by different agencies and prepare cloud best practices document.

2. BEST PRACTICES (OUR MANDATE)

The CGMS cloud expert working group submits a series of best practices based on lessons learned throughout our respective cloud journeys. This document reflects our collective current understanding and approaches to cloud while noting that this is an evolving area of implementation for all parties. The cloud is a disruptive technology that offers opportunities for advancement in science, data sharing and collaboration through the services it offers. Those advancements include but aren't limited to artificial intelligence and machine learning work, which will be drawing on cloud computing technologies.

This document focuses on best practices aligned to successes from CGMS agencies realised through demonstration and implementation of cloud capabilities. Beyond technical specifications, the best practices document focuses on organisational cloud adoption, including how to apply data standards, leverage new business processes and consider the disruption this will bring to the workforce.

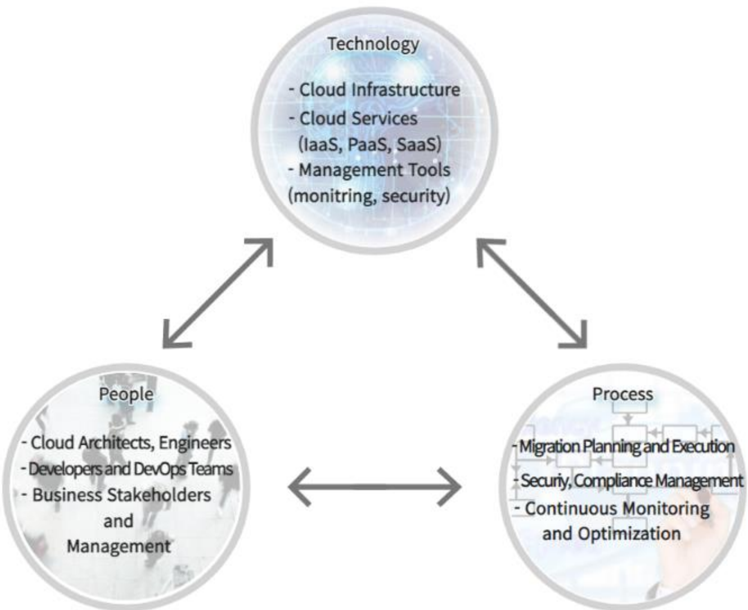

Figure 1. Describes this interrelation of technology, people and process which are all required for a successful cloud journey.

It is the CGMS cloud expert working groups intent to host annual workshops to continue to share our lessons learned and update this best practice document accordingly. This will continue to serve the needs of CGMS stakeholders as the document will be continuously refreshed to be reflective of cloud technical implementation, architectures, standards and cloud adoption elements to include business processes and people-oriented lessons learned.

The cloud journey and adoption

The cloud landscape is fundamentally more dynamic than the landscape of traditional data centers. This is in part due to characteristics of cloud resources, which tend to be less persistent and need higher redundancy than resources deployed in traditional data centers. Another contributing factor is the fact that cloud-native applications are often microservices rather than monoliths, with a higher number of individual components that might change independently of one another. This means that migrating into the cloud is a large change that requires careful coordination in a production environment, and following that migration, it is likely that smaller migrations will occur more frequently than in the previous environment. Therefore, adopting cloud technologies also means adopting a higher rate of evolution.

The frequency and nature of migrations and service adaptations are driven by evolving business and operational needs. As a result, cloud environments demand a more agile approach to data and service management, often resulting in more frequent updates and transitions. Managing this rate of change effectively is crucial. High change rates offer significant advantages for cloud systems, such as enhanced agility and responsiveness. It is equally important to minimise the burden on users, who otherwise might find themselves in a continuous cycle of adapting to these changes.

To ensure a smooth migration process, it is crucial to adopt a well-structured approach. This involves setting up systems in advance and effectively communicating deprecation timelines, and providing access to the new system in parallel to the old one, so that users can migrate without requiring coordination with other parties. A notable example is the transition of the World Meteorological Organization (WMO) Information System (WIS) to WIS2 in 2024. By providing a clear and staggered timeline, users are afforded the flexibility to migrate at a pace that aligns with their individual operational rhythms, avoiding the disruptions of a synchronised "big bang" switch-over. This approach should be applied to the initial cloud migration, as well as the numerous smaller migrations that will occur thereafter.

The key to successful cloud migration lies in strategic planning, taking into account what changes will affect whom, with what magnitude, and the costs for all stakeholders. This can be realised by frequent but manageable updates, and transparent communication with users. This approach mitigates the challenges of working in an environment with a high number of interconnected systems and supports seamless and continuous evolution for all users.

Understanding the need with thorough analysis

Cloud migration decisions should be based on mission requirements and a sound business case, and informed by a suitability analysis to confirm that a cloud-based alternative is viable. Careful and thorough analyses are critical to validate and select the most appropriate platform for the requirements. Hosted capabilities, strategy, and implications and benefits of refactoring/optimising legacy applications targeted for cloud migration should be considered.

Business case

Establishing a value chart that outlines an organisation's needs, values, and objectives is recommended as a best practice for cloud planning. At European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT), this led to the identification of a two-hemisphere cloud consisting of a Mission Critical Domain and a General Purpose Domain with an external cloud elasticity.

Cloud models

There are a three types of cloud models that organisations can chose from;

- Private Cloud: The infrastructure and cloud resources are used exclusively by one organisation and typically hosted on-site by the organisation or off-site through a third-party provider. The services and infrastructure are typically managed on a private network and responsibility for the maintenance is done by the organisation. The hardware and software are dedicated solely to the organisation's needs.
- Public Cloud: The infrastructure and cloud resources (storage, compute) are shared across multiple users (known as tenants) and typically managed by the public cloud provider. The infrastructure are typically maintained on a public network, accessed through the internet, and the responsibility for the maintenance is by the public cloud provider. The hardware and services are shared by all the tenants.
- Hybrid cloud: A combination of the private and public cloud.

There are advantages and challenges with the different models. Organisations should develop a cloud strategy that considers the business need, costs, consumer/users, and technical needs before finalising their cloud strategy. The best practices for cloud can be adopted across different cloud strategies and models. A few considerations that inform the cloud model and strategy are listed below.

Hybrid models

An integrated services portfolio brings the opportunity for cross-service integration by combining a variety of services such as data delivery, distribution, and data tailoring for cloud infrastructure. It is recommended that an Hybrid Cloud Management Software Orchestration be used to enable Multi-Domain Clouds which allows for coordination across other cloud networks. Hybrid models promote cloud-ready, on-premises object storage, elasticity and scalability. Additionally, since some public clouds and private clouds are architecturally incompatible, hybrid models allow for both systems to operate. Due to various architectural differences, choosing the right strategy will lead to the success of a hybrid cloud. The use of cloud technologies is the norm, it is however different across regions; the world will remain hybrid for a period of time. During the transition to a hybrid model, all assets need to be considered and choose what software and systems to migrate and what to leave on-premise. It is recommended to create 'landing zone' to accelerate migrations which has centralised network configurations, governance and policies, and understanding how you will deploy applications. Other concerns include data security, data copyright protection, and cloud migration for compatible adaptation when deploying a public cloud.

An enterprise cloud infrastructure is important for the implementation of ingest, processing, storage, science/development or distribution services. Organisational data should be ingested, created and stored centrally to minimise data routing and duplication of storage. Compute/processing and development of satellite data products can be done in a centralised service-based environment alongside the data to create a common cloud solution.

Governance

Effective governance for cloud shared services is necessary for cloud implementation and is should be established prior to opening a cloud-based service in the organisation. Cloud governance, with substantial organisational engagement, provides a strategic framework for cloud adoption and promotes consistency across implementations. Platform services under self-provisioning principles as well as data access services need to be considered in governance mechanisms. Governance is supported by suitable monitoring mechanisms such as Key Performance Indicators (KPIs), accounting, and any applicable (data) policies the organisation might have. This should occur at both the strategic level (e.g., policy, the selection of cloud vendors and brokerage models) and operational level (i.e., running the cloud), and addresses who can do what, when, and to which extent within the cloud. Effective governance also provides management mechanisms and structures to ensure that the use of on-demand cloud-based computing services is conducted within budgetary constraints. These structures should empower developers as much as possible to have necessary services and resources as they need for their tasks when they need that with as little bureaucracy as possible.

Governance encourages cooperative behaviors, promotes best practices, establishes performance metrics, facilitates compliance, and centralises communication and coordination of cloud activities. These aspects are essential service design drivers and must be addressed at the beginning stages of cloud adoption, and have proven to be difficult to correct afterwards.

Another area of consideration is having a proper account structure, policies, and a service catalogue. Policy automation are described in more detail in below. The account structure helps with costing in a service model and detailed below under Financials.

The consumer of cloud services

A thorough understanding of the business customer needs and the target audience should be done prior to selecting the service provider. The analysis should include at least geographical aspect, scalability issues, required downstream services / required features of the cloud platform, availability and continuity requirements, cost-model, whether users are more willing to utilize large public cloud providers, or limited access local cloud implementations, and security and data privacy issues.

Scalability of service must be considered when choosing a cloud service provider. If the data is to be shared with the global community, the option of automated replication of products in different regions is another consideration.

Redundancy and business continuity

Public cloud providers operate in several regions. This enables the organisation to develop a redundancy and business continuity plan that can avail of a multi-regions, multi-cloud environment. This enables organisations to execute on their plan for business continuity, availability, continuous operations and be part of the Disaster Recovery (DR) strategy. Infrastructure as Code (IaC), discussed below plays a large role in enabling organisations execute on their business continuity plan.

Financial strategy

Agencies must adhere to their government rules and regulations to manage competition between cloud vendors with a view towards adopting changes as necessary. Organisations can have several different funding codes and allocations representing different offices within the organisation. An enterprise cloud framework approach that each office leverage can assist in financial management of service usage and tracking. Implementation of a tagging schema against cloud resources can allow organisations to assign funding-specific tags to data and services to manage across multiple budget schemes. This also ensures compliance to fund tracking while using a common enterprise service. can Cloud transformation should result in changes not only in the technical departments but rather across the entire organisation. Cross-functional teams in engineering, finance, operations, etc should work together to achieve better and faster result, gaining at the same time more financial control and predictability of the cloud usage. Typically, this is done by introducing cloud FinOps practices in the organisation. FinOps would be defined as an operational framework to help with decision making through collaboration from Finance, Engineering and Business Operations. The goal is for organisations to have a better understanding of cloud costs to help execute on their overall strategy.

There's another option that each team create their own account which is attached to central billing account and/or contract. This would empower the teams even more but requires team-based budgeting scheme in the organisations.

Additionally, when working with a commercial cloud vendor to reduce "vendor lock-in," leverage services that are equivalent or have capabilities of all potential cloud service providers. Agencies may also consider using multi-clouds to reduce vendor lock, especially if they don’t want their data to be lost; this defers potential risk. Another option is establishing a long-term contract for a fixed price considering lock-in as part of the lifecycle for the application and align the migration of the application to the new vendor with the lifecycle of the application.

The following financial aspects must be managed, considered, and negotiated if possible, when implementing cloud services:

- Egress costs
- Capital assets vs operating costs
- Discounts from cloud vendor
- Vendor lock-in
- Licensing of software based on capacity
- Ownership of the infrastructure as code (data/code)

Workforce and users

Empowering the workforce for cloud-readiness is key for stakeholder buy-in. Cloud impacts all facets of the workforce, from scientists and engineers to acquisitions to program management. The workforce must have the skills, competencies, authority and information needed to apply cloud technologies to their day-to-day job. Preparing and re-training the workforce will enable new process development and promote innovation, which in turn could increase top talent recruitment. This should be followed by a plan that addresses roles, responsibilities, and workforce impacts, as cloud implementation does not just affect IT professionals. The plans will have the opportunity to address skill gaps, training schedules, and curriculum that will be required for the entire workforce (e.g., budget analyst, acquisitions). The workforce requires a careful analysis and use of automated tools to ensure the enterprise is adequately able to meet its growing mission.

Correctly governed cloud technologies can also have a strong empowering effect to developers when they get correct resources and services when needed. Being able to self-manage the resources removes unnecessary obstacles in the development process and makes agile elaborating best possible solution and infrastructure requirements possible. Being able to use ready downstream services enables developers to focus on the domain-specific task.

Most importantly, cloud services must retrieve feedback from the users and workforce so they can leverage the cloud beyond the data. Their feedback will provide valuable information to improve their cloud services.

Migration strategies

Some of the different strategy options when migrating to the cloud are included in the table below:

7R	Description	Pro	Cons
Rehosting (Lift -and-Shift)	Lift-and-Shift applications to the cloud using virtualisation. This involves moving applications without changes. Moving applications from the local environment to the cloud using tools like Server Migration Seraebabvice (SMS) from AWS or manual procedures	Quick and simple, minimal changes to applications, does not require a deep understanding of the cloud environment	Might not take full advantage of cloud-native features, might lead to higher costs if applications are not optimised for the cloud
Replatforming (lift-and-Reshape)	Create a completely new platform to offer the benefits of the lift-and-shift approach, while also capitalising on more of the performance opportunities the cloud can offer. A few cloud optimisations are made to achieve some tangible benefit but aren't changing the core architecture of the application	Requires only minor optimisations, can provide better performance and cost savings than rehosting	Might require more time and effort than simple rehosting, still does not take full advantage of cloud-native features
Refactoring (R ewriting/De-	Completely redesign code to ensure it makes the most of the cloud. This is ideal way to migrate to the cloud. This involves	Allows for full use of cloud-native features, can lead to better	Requires the most effort and resources, can lead

7R	Description	Pro	Cons
Coupling Apps, Re-architect)	reimagining how the application is architected and developed using cloud-native features. This is often driven by a strong business need to add features, scale, or performance that would otherwise be difficult to achieve in the application's existing environment	performance, scalability, and cost savings in the long term	to potential risks and downtime if not properly managed
Reimagine	This involves transformational activities that include going beyond just moving to the cloud, such as rethinking IT roles and policies or modernising your organisation's IT development culture and processes (e.g., adopting DevOps).	Can lead to significant innovation and improvements in business processes, allows for full use of advanced cloud technologies such as AI and machine learning	Requires significant effort and change management, might involve risks and challenges associated with new technologies
Repurchase (Replace – Drop and Shop)	Purchase alternative services from your vendors that are designed specifically for the cloud. This is a quick way to get the full benefits of cloud. This could involve moving application from on-premises software to its Software as a Service (SaaS) version	Can improve performance and reduce maintenance effort by moving to a different product, simplifies the migration process	Can lead to compatibility issues with existing systems and data, might require significant changes to operations and workflows.
Retain	You might do nothing to certain applications if they're not ideal for a cloud environment. However, retaining applications may be a temporary decision while you're working on one of the other paths	No immediate cost or effort required for migration, allows time for further assessment and planning	Does not provide any immediate benefits of the cloud, maintenance costs for on-premise systems continue
Retire	If after reviewing your applications, you find ones that are no longer needed, these can be turned off. This can save costs for the company	Reduces costs and effort by eliminating unneeded applications	Might lead to loss of functionality if not properly assessed, can disrupt users and workflows if not managed properly

Once a migration option or a combination of options has been chosen, it is best to define a migration strategy for the migrations.

Migration Execution

Organisations will typically undergo several phases of their migration journey. Depending on the scale of their migrations, they should consider developing a migration factory design. Key considerations when developing a migration factory design include pre- and post- migration activities, testing, training, and retiring the legacy applications that might be on-prem. Considerations for record retention need to be done throughout the process. A cloud costing model should be developed that can be leveraged to determine costs each migration candidate application or suite of applications.

Cloud pilots and demonstrations

Pilots are recommended when beginning cloud services implementation. The duration of the pilot phase should be dependent on pilot in question, taking into account e.g. the critical need of the application. Small scale environments for applications training has been proven to be very successful and can lead the path to a bigger implementation of the cloud. Pilots allow ongoing implementation and demonstrations in a simple, safe to deploy package, and instils trust in cloud technology. Additionally, a pathfinder approach can help small projects find specific cloud data services while explore new cloud technologies. Moreover, combined with the right data, cloud pilots and demonstrations can develop use cases to leverage a public cloud and/or commercial services.

3. BEST PRACTICE FOR CLOUD ARCHITECTURE AND IMPLEMENTATION

Planning for cloud architecture

When migrating the existing services running on traditional on-premise IT to cloud two approaches are available.

Lift & shift approach is usually quick and simple way to mimic already existing system or service by simply recreating same software architecture 1:1 in the cloud by deploying a software components on the virtual environments. These are usually applications that have a steady load, don't require any dynamic up- or down-scaling and very static in terms of configuration and software updates. Cloud native approach on the contrary requires drastic changes to the application design. Typically, monolithic application needs to be broken down into the set of various micro services, ideally leveraging the existing portfolio of cloud provider. This usually leads to a major re-design of the application around the capabilities of the chosen cloud provider, which on one hand allows more reliable, performant and cost-efficient application, but on the other hand may introduce vendor locking. Another hybrid approach can be taken when migrating applications to cloud, whereby migration to cloud-native takes into consideration common denominator of services among all public providers of your choice. This will reduce the risk of vendor-locking and allows migration to cloud still using cloud-native approach.

It is also important to design the architecture to be cloud-native in practice; this will enable scalability, containerisation, object storage, and affinity groups which will benefit processing features.

A best practice for setting up cloud architecture is to create two areas for operations and research and development. The operations (routine production) area is used to build an automated approach to products that are done routinely. It enables a process-specific system that is high-performance computing and scalable from one clear access point. This will overall bring people to the data and increase the value of the data through automation and routine production. Additionally, in order to enable all heavy workloads such as modelling behaviours, your architecture needs to be able to accommodate different formats including leveraging higher complex capabilities. The cloud is a platform to enable this and provides the benefit of enabling high data analytics.

The research and development (R&D) area is for scientific developers, cooperative institutes, and partners to collaborate for the coding and testing of data. This will allow R&D algorithms to deploy faster using the dev ops approach. It brings the developers to the data to work in a virtual environment to reduce copying of data.

Ingest

Ingestion aka data ingestion is the process of transporting data from one data source to another. Data ingestion enables organisations to make sense of the complexity of data by ingesting into the cloud and into one of several technologies such as data lake, data mart and warehouses. Depending on the need, organisations can chose from different ingestion strategies including real-time, batch or a combination. Data ingestion tools can address different features and capabilities including format, frequency, size and privacy needs of the organisation.

Storage

A best practice for storing data in the cloud is to leverage object storage architectures available as a service in all cloud providers as well as on-premises.

To best use object storage, it is recommended to make full use of its main advantages over traditional methods. Data stored should be enriched with meaningful metadata that will simplify data handling for users and applications. Built-in data lifecycle management should be used to automatically move data between different storage tiers based on frequency of access, latency and throughput needs. Leveraging the standard S3 API across multiple cloud providers is a benefit for users and application developers, but providing access to data directly via S3 APIs should be considered only in exceptional cases.

In the absence of widely adopted data formats and data discovery standards, an abstraction layer in front of the stored data is necessary. The harmonised data access layer will simplify data usage for developers and end users while at the same time allowing the implementation of additional logic not possible directly in S3. Examples of needed functionality would be related to data policy and data licensing.

Another strategy to reduce storage cost whilst not affecting performance is to enable tiering of data. With automated tiering, active data can be stored and tiered separate from more inactive data. With this approach, active data would remain in high-performance storage, while inactive data is tiered to low-cost storage.

Metadata catalogue

Metadata catalogue is a collection of data about the data. It includes key attributes that one wants to collect about the data and can include the data owner, source, origin, sizing and other attributes about the data set. It can serve to be a repository of the structure, definitions and quality of data. A rich metadata catalogue can be a great tool for filtering and searching the data. The use of a common catalogue and metadata inventory will aid the ingest, storage, compute, and dissemination in the cloud. Metadata inventory and cataloguing support all data storage in the cloud environment, including data that is ingested into the system, new products that are generated, and external access to the stored data. Metadata provides the foundation and is the key to developing a common way to store, search and retrieve data across the framework. Metadata also enables the data to be searchable and increases data discovery. The cloud can optimise how your data is being stored and disseminated.

Processing

The cloud enables fast and efficient provisioning of computational resources. Utilisation of virtual systems for database and storage services in the cloud is a processing best practice. Virtualisation helps in improving scalability, reducing cost, promoting optimal utilisation of resources and helps making system highly available by running multiple operating systems on a single hardware, dynamic allocation/re-allocation of resources based on requirements and on-the fly migration from one system to another.

Cloud processing of data provides high performance computing, a one stop service for applications and publishing. Data-proximate cloud resources can be implemented in hybrid environments. Automation on every level should be the objective for every cloud system since it will enable a virtual infrastructure, architecture deployment, operations, and observability (monitoring, logging, alerting), giving the cloud sponsor a financially sound and security perspective of their IT system.

Data

The needs and requirements of users can vary widely in terms of reliability, timeliness, and availability, as well as in terms of their need for compatibility with legacy systems. Recognising this, members of the CGMS Working Group have developed distinct approaches to cater to different user demands: Data distribution for real-time data applications, and pull-based data access. Each of these styles of data access are implement using different systems and services and are described in greater detail below. The bifurcation of approaches represents a thoughtful and cost-optimized response to the diverse needs of the user community, ensuring that whether it's the immediacy of real-time data or the flexibility of pull-based access, the right solution is available. As we delve into the details of these systems, we'll explore how they have been tailored to meet specific needs and how they continue to evolve in a rapidly changing technological landscape.

Data access and delivery/dissemination

Delivering data, and especially near real time data, is challenging in the context of exponentially increasing meteorological product sizes. Different strategies can be employed to counter the problem. In the context of cloud storage, one part of the solution could be the usage of content distribution networks (CDN). CDNs can allow full control and ownership of the data, while at the same time delivering the right data close to destination hotspots. CDNs work best in a pull model where the client pulls the data from a location either automatically or manually. This is most suitable when only certain, but not all, products are needed, so that it does not make sense to use the entire set of available data.

Implementing a pull model for data distribution usually requires a secondary technical solution to alert users to the availability of new data products. Regardless of whether the data is stored in one or more data lakes or if CDNs are used or not, the producers of data need to effectively communicate with the consumers of it. In modern cloud architecture, applications and data are decoupled into smaller, independent building blocks. Messaging provides instant event notifications for cloud distributed systems. It supports scalable and reliable communication between independent software modules. If a publish-subscribe model is used in the architecture of data distribution, it is recommended to use available messaging products and services that support it, rather than building your own.

For instances where users want access to all incoming data, a subscription model can be used in order to stream the data directly to consumers as soon as it is available. Because the user doesn't need to decide whether to download a specific product or not, the overhead of alerting them of data availability and then dispatching a request for that data from the user's side can be saved.

The challenges of data access in the context of weather product reprocessing, climate science and Artificial Intelligence (AI)/ Machine Learning (ML) model training is different compared to data dissemination. In such use cases, having fast access to large amounts of historical data is the focus. It can be useful to put such data into a data lake, so that diverse products can be accessed from the same system. The usual S3 data tiering strategies are counterproductive as most of the time the data needed is not the most recent data produced. A different data tiering model when should be used when designing the data architecture. The data tiering should be progressively smaller in size and exponentially faster in throughput and latency as it is moved closer to the processing function. It is important for the end processing function to be able to proactively move data through the storage tiers before it is needed. Another important feature of the data should be the possibility to properly discover and access the substructure of products such that the amount of data moving through the data tiers can be minimised without the need to use extra computational cycles to modify it.

Data life cycle

The life cycle of data that is tiered in CDNs and/or fast storage (in the context of data dissemination) should be defined. Apart from the central location of the data, replicas should be removed when data expires, or all registered consumers have received it. The same is the case for tiering data in the context of massive data processing where higher performing storage tiers should maintain the data only as long as there is a high chance of a future need to access it. For data processing, the lifecycle is driven primarily by the cost trade-off between re-accessing the data from the data lakes and the overall cost of storing the data in higher performance storage tiers.

Real-time data applications

For users with the most stringent requirements, such as those involved in numerical weather prediction, real-time services are paramount. Multiple agencies have long-standing systems that push data out to users with guaranteed timeliness, reliability, speed, and availability. These systems, being the oldest, have proven their worth in meeting the critical demands of real-time applications. However, it's worth noting that the data pushed out with these systems isn't persistent; users must store it if they wish to keep it.

From the early days of the Internet, agencies have used commercial satellites to disseminate data to users with a dedicated earth stations. As terrestrial networks became more reliable, trade offs revealed found that terrestrial multicast in the same style of satellite broadcast scaled excellently in terms of cost, outperforming CDNs. This led to the use of multicast inside multicast-enabled networks, specifically national research and education networks (NRENs), to distribute large volumes of data quickly. Both satellite dissemination and terrestrial remain very attractive from the perspective of cost when there is a large number of users requiring high data volumes with stringent timeliness constraints, and the majority of users require a large portion of the available real-time data.

As user communities have grown to encompass users without access to NRENs and satellite ground stations, innovations like automatic multicast tunnelling (AMT) now enable data distribution to the wider web and various clouds without the need for specialised hardware. This method is now used to serve users on every continent except Antarctica, providing excellent service through a cost-effective platform that accommodates both satellite and pure software solutions.

The WMO is also contributing to this landscape by building the WMO Information System 2 (WIS2). Utilising pub-sub technology, WIS2 announces the availability of data, which is then pulled into a cache. Though it lacks guaranteed network bandwidth, the small volume of subscription data allows it to function similarly to real-time applications, fulfilling many of the same demands for operational weather centres without the need for a proprietary network.

When resources are not available to setup such highly optimised services, or where the preconditions for these optimised multicast services to be attractive are not yet, cloud operator’s services can be leveraged to provide real-time access to data products via HTTP or FTP, possibly supported by a CDN layer.

When considering the use of cloud providers' services, there are several factors to take into account. Utilising a cloud provider's services can offer faster access, CDN-like distribution, and sometimes lower costs. However, dependence on a single cloud provider can limit flexibility and options due to vendor lock-in. To mitigate this, agencies might maintain independence through traditional internet connectivity or by utilising multiple clouds, catering to a broader user base.

Scalable storage solutions are crucial for efficient data distribution, accommodating both structured and unstructured data, including time-sensitive information. Subscription systems ensure that data is automatically delivered to users according to their needs. Platforms should be provided so that users can select specific data sets in line with their needs. Once available, the user can be notified or the data transmitted to the users' preferred storage, ensuring seamless and flexible access.

Access

For users seeking access to non-real-time applications, complementary services can help users tailor data offerings to their specific needs. Visualisation services are excellent for exploring data products and can also fast-track downstream applications if they make these visualisations available via widely adopted APIs. Discovery, download, and tailoring services for historical data are also helpful in prototyping cloud applications that consume satellite data, in cases where only small amounts of data are needed, or where long time series are the priority. Throughout, such services should be offered with graphical user interfaces (GUIs) and application programming interfaces (APIs) so that humans and machines can consume the services. Agencies would also do well to provide computing environments that are located in the same network as the data services so that users can implement data-proximate workflows in the same environment as the data offerings. This can also make it easier for agencies to offer specialised tools to assist users in taking advantage of their desired data products.

Secure data access in the cloud involves only storing approved, non-sensitive data for the user community. After uploading the data, the required permissions should be checked and verified in an automated mode. There should not be any global data shares uploaded on the cloud and a periodic audit of the cloud account including monitoring of access logs should occur. When making data available, it’s crucial to test and document the most typical use cases employing generally available and used tools and libraries. Additionally, data upload on the cloud should employ optimal network bandwidth with a configured firewall to ensure timely upload. There should not be any directed network connectivity of the operational systems in the cloud, and data should be routed through designated gateway/intermediate systems. Interfaces should be simplified.

Security

Historically, if confidentiality, availability, or integrity of data needs a high security designation, the entire system would need to be classified and built to that high standard. With the cloud tailoring security features, agencies can adjust the service and tool security controls to meet the intent of the high security system on an individual basis. Agency compliance requirements must be met in addition to understanding a shared security model across the agency. Security considerations for the cloud are:

- Location of data storage
- Educating and training staff on security requirements
- Data encryption
- Securing data endpoints
- Targeted monitoring
- Audits and penetration testing (and use of 3rd party evaluations to assist)
- Intelligent Security logs
- Multifactor authentication
- Establishing security notifications and alerts
- Scanning and patching containers
- Enhanced security logging; reviewing costing information can convey abnormal behavior in the system.
- Identity management solutions and access: determining which users have access to change accessibility rights. Automated access role enforcement is key in automated tools as well.

Infrastructure as code

Infrastructure-as-code (IaC) is the automated management and provisioning of infrastructure. This includes servers, security settings, cloud resource deployment, version upgrades, backup and data recovery, and security administration. Benefits include:

1. Agility and Cost Reduction: What once took days or weeks to setup and configure, infrastructure can now be done in minutes.
2. Simplicity: With the use of automation, dozens of manual steps can now be automated as part of the DevSecOps pipeline and will be simpler to propagate from one environment to another.
3. Repeatability: IaC can improve consistency and repeatability of the infrastructure, helps to achieve same results during testing, help avoiding configuration mistakes.
4. Scalability: IaC enables organisations to address increased demand when needed and resources can be dynamically added or removed.
5. Reliability: With the infrastructure setup automatically, there is an increased reliability that the infrastructure will be setup correctly and have less human errors.
6. Reduce vendor lock-in: As part of a vendor lock-in, organisations have more flexibility to move and/or migrate their infrastructure to another cloud vendor if their servers have been built thru' IaC.

Service orchestration and automation platform

SOAP enables the design and development of business services through a combination of workflow orchestration, run book automation along with the provisioning of resources. With the adoption of SOAP, organisations can automate use cases that deliver and extend into data pipelines, cloud-native infrastructure and application architectures.

With SOAP organisation can effectively and efficiently manage data pipelines, the orchestration of application workflows, and the provision of resources. Some of the benefits include:

- Manual steps can be automated through event-driven automation.
- The scheduling, monitoring, visibility, and alerting of resources
- Resources can be provisioned for both on-premise and cloud including storage, network and compute
- Workflows can be better managed and automated

Agile and DevOps methodologies

When implementing cloud services, it is strongly encouraged to leverage Agile and DevOps Methodologies resulting in increased efficiencies. The cloud offers an opportunity to rapidly provision IT resources, promote scientific innovation and build automation of routine processes. These opportunities help reduce research to operations transition time and produce organisational cost savings, but can only be realised if institutional processes are realigned for cloud implementation. While cloud can be implemented through waterfall processes, migrating to an agile methodology is key as it aligns better with the shorter timescales and flexibility the technology provides. One successful approach that preserves separation of duties entails moving toward a DevOps practice where strict handovers and engineering reviews are homogenised into one team responsible for iterative development and operational deployment. This ensures that all development activities are can be released as rapidly as needed and possible (depending on the degree of automation) and incrementally into operations. This approach brings together scientists, systems engineers, operational monitors, and security and cloud experts into one algorithm-specific team. As the cloud is being implemented using these approaches, ensure to update documentation practices and provide to the cloud workforce for deployments, configurations, etc.

Considerations for cloud native data formats

To take full advantage of cloud capabilities, agencies should consider transforming their data into cloud native formats. These formats have been designed specifically with use in the cloud in mind, and typically prioritises scalability of access, asynchronous read/write capabilities, streaming (sometimes via API), and metadata access without requiring access to the data itself. This allows multiple applications to utilize these data simultaneously. It also allows them to discover and access only the portions of data products that are relevant at a given point in time, rather than having to download and process the entire product, which can be very large.

Cloud native data formats can offer several advantages in a cloud environment. However, it is not always the case that data is produced in these formats. There are various reasons for this - legacy processing chains, historical data, streaming strategies, etc. are some but not all of the valid reasons for having data in a format that is not specifically optimised for the cloud. Transforming large volumes of such data into cloud formats can be costly, and may also impact users. Agencies should carefully analyse the costs versus benefits of changing their production strategies or transforming their existing data before pivoting to new data formats, as in some instances the costs will outweigh the benefits.

Implementation at the enterprise, global collaboration through the cloud

If at all possible, global collaboration is the best solution. As agencies move from an agency-centric system to a collaborative earth observation model including commercial satellites and international partners. The volume of data and products have increased significantly, and have transitioned from a point-to-point connection to sharing data on the internet. The volume of data has gotten so big that point-to-point sharing is overloaded, and the cloud presents a solution. "Best of breed code" from multiple agencies can help agencies utilize the most modern code, allows for local customisation of code, mutual access for public infrastructure, and facilitate remote validation work.

4. CONCLUSION

Cloud computing technologies offer organisations unprecedented power and speed. This applies not only to the capability to rapidly and flexibly allocate infrastructure for new initiatives, or to scale that infrastructure in line with business needs, but also in the ability to better control application deployments and roll out updates more safely and speedily. This comes at the cost of complexity due to new technologies and working practices.

The advantages of cloud computing can only truly be harnessed if decision structures within the organisations utilising clouds are capable of leveraging them. The pace of change in clouds can be high, and the environment can be more complex. This means that traditional command-and-control authorisation structures, often exemplified by phase gate models, will not unlock the benefits of the cloud. Instead, a guard rail approach should be used with decision making authority moved as close to the work as is possible and safe, with clear boundaries of responsibility and authority defined. In a best case these boundaries should be reflected in technical capabilities so that users in the cloud can experiment safely without fear of overstepping their bounds. Because this delegates significantly more autonomy to people executing the work than before, it may also require updates to the financial procedures of the organisations, especially when working completely or partially in public clouds.

Clouds can also be utilized across agencies as a common ground in which collaborative efforts can take place. This means that multiple agencies can collocate data and processing, and leverage a common set of tools, with much less effort than previously was required. By using these capabilities, it is possible that inter-agency collaboration is greatly improved, augmenting the capabilities of the global community to deliver greater benefit at a higher pace than ever before.

Cloud technologies represent a paradigm shift in computing. Due to the critical nature of the work of CGMS members, it is essential that the ramifications of this paradigm shift are clearly understood and their impacts analysed. Only in this way can the substantial benefits of the cloud be unlocked for agencies and their users without endangering the current service portfolio. When executed carefully and methodically, there is potential to greatly potentiate the benefits that agencies can provide to their employees, users, and the global community at large.

Appendix A: Acknowledgements

The following CGMS WGIV Cloud Experts sub-group members contributed to the current version of this document:

- CMA, Xu Zhe
- EUMETSAT, Klaus-Peter Renner
- EUMETSAT, Borys Saulyak,
- EUMETSAT, Ismael Cuadrado Cordero
- EUMETSAT, Roope Tervo
- EUMETSAT, Mike Grant
- EUMETSAT, Daniel Lee
- ISRO, Nitant Dube
- JMA, Kotaro Bessho,
- JMA, Masaya Takahashi
- JMA, Yasutaka Hokase
- KMA, Sung-rae Chung
- KMA, Dae-gyeom Jeon
- NOAA, Kathryn Shontz
- NOAA, Kumar Machado
- NOAA, Jessica Wieman, CGMS Secretariat
- NOAA, Maria Grady, CGMS Secretariat