WHITE PAPER

Cloud Computing

Raising Geospatial Technology to the Cloud: Intergraph® Strategy for Leveraging Cloud-based Resources

INTERGRAPH®

Contents

1. Introduction 1

2. Cloud Computing 2

2.1. Clustering 2

2.2. Connectivity 2

2.3. Abstraction 2

2.4. Convention 3

2.5. Culture 3

2.6. Evolution in Datacenter 3

2.6.1 Traditional Datacenter 3

2.6.2 Virtualized Datacenter 4

2.6.3 Cloud Computing Enabled Datacenter 4

2.7. Cloud Computing Definition 4

3. Cloud Architecture and Business Models 5

3.1. Public Clouds 5

3.2. Private Clouds 5

3.3. Cloud Infrastructure and IaaS 6

3.4. Cloud Platform and PaaS 6

3.5. Cloud Application and SaaS 6

4. Impact of Cloud Computing on the Geospatial Landscape 7

4.1. Publishing Spatial Data on the Cloud 7

4.2. Processing Spatial Data in the Cloud/High-Performance Computing 7

4.3. Backend Services for Mobile Applications 8

4.4. Geospatial Complex Event Processing in the Cloud 8

4.5. Inhibitors to Cloud-based Solutions 8

5. Intergraph's Cloud Computing Strategy 9

5.1. Application Architecture for Cloud Computing 9

5.2. Certifying on Existing Cloud Platforms 10

5.3. Cloud Management and Licensing Tools for Geospatial Solutions 10

5.4. Cloud Services Consumption and Integration 10

6. Conclusion 12

1. Introduction

Cloud computing is one of the most talked about trends in information technology today and it has already significantly impacted the IT landscape. Major players such as Microsoft®, Google, and Amazon are making massive investments in cloud computing infrastructure as they make serious financial bets that cloud computing will continue to gain market share.

Is cloud computing a viable deployment model for Intergraph® applications? Does Intergraph believe that cloud computing concepts will be embraced by our customers? In particular, geospatial processing has some unique aspects when compared with general IT, such as the requirement for large volumes of data and specialized tools for visualization. How do these factors affect geospatial processing's suitability for a cloud computing approach? To answer these questions, this white paper examines some of the background and factors enabling the cloud computing phenomena. It also discusses cloud computing's applicability for Intergraph's solutions and geospatial processing specifically, and sets out Intergraph's strategy and approach for leveraging cloud-based resources.

2. Cloud Computing

One can argue cloud computing first began during the days of timesharing on mainframe computers in the 1960s and ’70s. Cloud computing has evolved into what it is today through the influence of various concepts and technologies, such as grid computing, utility computing, software as a service (SaaS), application service provider (ASP), virtualization, multi-tenancy, etc. Ultimately, cloud computing is a culmination of experiences and advances in areas such as clustering, connectivity, abstraction, convention, and culture. The following section reviews some of these factors in more detail.

2.1. Clustering

Key to the cloud model is the reliable provision of computing cycles. How much computational power can be packed into a square meter with minimized cost? A few decades ago, integrated circuits revolutionized the production and performance of electronic equipment. The revolution from discrete electronic components to integrated, cheap chips led to enormous improvements in the efficiency of computing operations and allowed logic to be integrated in the range of devices we see today. This same densification has continued in the computing world with multiprocessor and multicore processors and continues further with clustered racks of blade servers in data centers and now with integrated containers of computing equipment or modular prefabricated data centers to further enhance performance, efficiency, and automation. This clustering effect provides compute cycles with a critical mass and a cost per compute cycle point that traditional IT deployments cannot match.

2.2. Connectivity

Where is the most efficient place to perform the computing task? Connectivity enables effective deployment of computational power on all levels in an application architecture. Each discrete element has to be connected in some way with the rest of the system. Fast, reliable, cheap interconnectivity is the cornerstone of various communications types – whether on chip, a single computer unit, within or outside the data center. The Internet has revolutionized the way we utilize information, access application functionality, and perform social interactions. Ubiquitous network connections enable access to cloud computing from a variety of applications, devices, and workflows. Mature connectivity allows us to delegate tasks to specialized applications in the cloud that can use the full power and elasticity of the cloud computing infrastructure.

2.3. Abstraction

Modern cloud architectures simplify access to computing resources. To what degree do we need to deal with IT infrastructure details to solve business problems or execute operational tasks? Abstraction enables flexibility and facilitates the right scale of resources to operate on a given problem. In the software industry, abstraction has been achieved by various technologies and architectural designs, most notably and recently service oriented architecture (SOA) and virtualization. SOA is an architectural style to abstract business functions away from the concrete software provider and a style for scalable, distributed computing, while virtualization enables users to fully decouple runtime applications or services from the dependency on the underlying hardware infrastructure. Cloud computing typically uses virtualization internally for scaling, elasticity, and reliability, while applications delivered to the end users use SOA.

2.4. Convention

What is the right ratio of uniformity versus specialization? Conventions or standardization helps to decrease unnecessary variability and improve efficiency in manufacturing and integration. This leads to cheaper components and to commoditization of the base IT elements. Cloud computing leverages commodity hardware and standards in communications and application delivery. This enables economical operations to run in the cloud on the provider side, as well as easy access from the user or software vendor side. A typical cloud environment runs on cheap but highly redundant hardware infrastructure that enables high reliability. More specialization is occurring on the software side.

2.5. Culture

Are we comfortable with our processing and data existing in the cloud? During the past decade we have been conditioned to take ubiquitous access to the network for granted. Moreover, Internet, social media, Web applications, Web mail, Web storage, peer-to-peer networks, videoconferencing, and many other styles of communication have become popular and widely accepted by consumers. Microsoft, Google, Amazon, and others use cloud computing to deliver these experiences. Organizations are now beginning to take advantage of these cloud computing concepts to deliver on their mandates.

2.6. Evolution in Datacenter

Figure 1 illustrates the cloud computing evolution in the data center with respect to resource utilization, agility, and ease of operations.

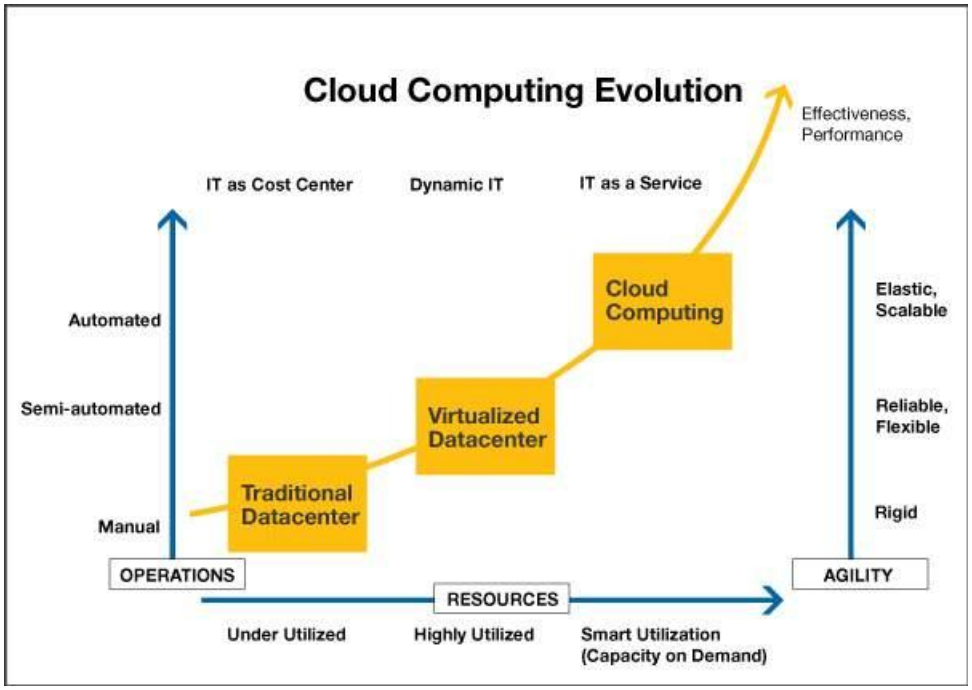

Figure 1: This diagram shows the evolution of cloud computing in the data center with respect to resource utilization, agility, and ease of operations.

2.6.1 Traditional Datacenter

Traditional datacenters, similar to electronic components in the early stages of computing, typically feature individual servers that are isolated, underutilized, and expensive to maintain. This leads to high

power consumption, physical space constraints, and a high cost to provision the full data center, especially for mission-critical systems that require redundancy and failover protection.

2.6.2 Virtualized Datacenter

Virtualization, similar to integrated circuits on silicon wafer, brings the concepts of integration and densification to the data center. Virtualization provides a hypervisor that acts as an underlying platform on a physical machine to integrate various operating systems and the related software stack. While hypervisors were originally introduced on mainframes, modern hypervisors enable the use of the same virtualization technology on commodity hardware. This minimizes the number of physical computers needed, increases efficiency, decreases the overall weight of the datacenter, and provides many other useful features, such as high availability, security, and scalability.

2.6.3 Cloud Computing Enabled Datacenter

Cloud computing progresses the integration and densification so that the whole physical infrastructure appears as a single coherent and extremely efficient environment where deployed software instances can fluctuate according to the operational demands. A cloud computing datacenter acts as a large computing appliance integrating many inexpensive computing units (usually standardized as blades) that include efficient support for virtualization. Virtualization management tools, or cloud operating systems, ensure automated operations, efficient provisioning/deprovisioning of physical resources, self-healing, load balancing, reliability, and security. Cloud-based datacenters hide internal complexity, significantly reduce maintenance required from IT staff, and offer a simple user interface.

2.7. Cloud Computing Definition

Cloud computing is a next-generation software (application and services) hosting technology that can be owned and operated by an organization (the private or internal cloud) or by independent provider (public cloud).

Cloud computing uses the latest innovations and experiences across computing technologies to maximize efficiency of operations, while minimizing the cost of computing components, electrical power, space, and time required to provision additional resources. Efficiency is reflected in typically using commodity components and uniformity with imposed conventions and standards that might be delivered as a “cloud computing appliance” that ties together pre-integrated and preconfigured self-healing computing systems. This dramatically simplifies maintenance, reduces the need for on-site visits, and provides turn-key solutions for provisioning additional cloud infrastructure resources. All interaction with the cloud infrastructure is done through Web services or self-service Web portals for unified management. Clouds can host heterogeneous applications, provide multi-tenancy, and precisely measure resource usage per application.

Cloud computing uses highly parallelized virtualized computing infrastructure with fully automated operations to provide failure resiliency, load balancing, and agile elasticity in scaling up and down without assistance from IT personnel. This makes it possible to dynamically and on demand allocate or de-allocate new computing resources for running business applications.

3. Cloud Architecture and Business Models

Cloud computing architecture consists of several layers that offer various levels of services and capabilities, and enables various business models, deployment scenarios, or usage. To help understand the relationships of these layers, Figure 2 illustrates the typical cloud computing architecture in relation to public and private clouds.

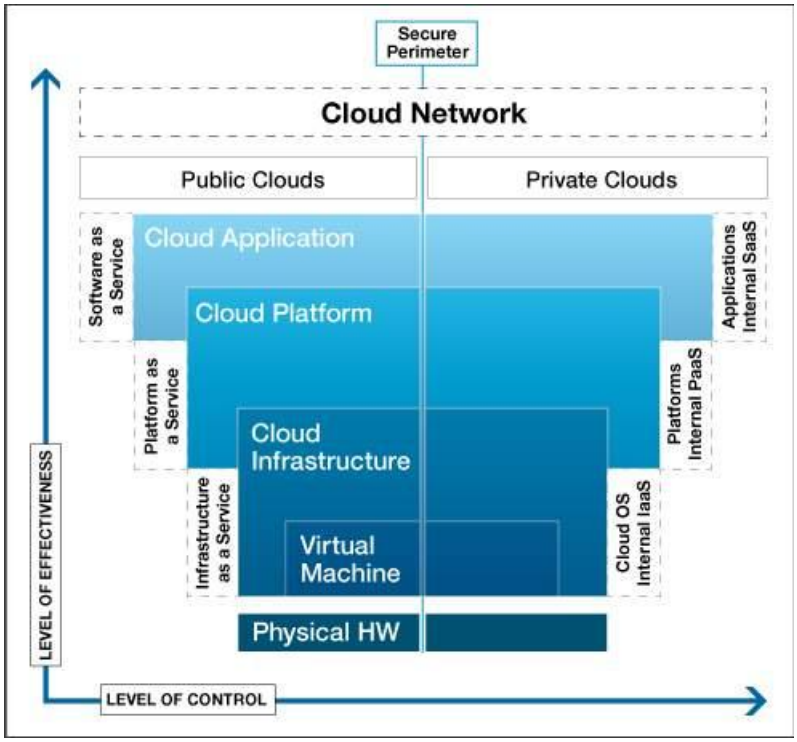

Figure 2: This diagram illustrates typical cloud computing architecture in relation to public and private clouds and level of control and effectiveness.

3.1. Public Clouds

Public clouds are publicly available, multi-tenant pools of computing resources from an independent cloud provider that build on economies of scale paradigm. IT resources are provided "as a service," fully automated, based on a service level agreement and charged to customers per actual use. Public cloud providers specialize in the efficiency of building and operating their data centers through a massive aggregation of a large spectrum of demands for IT capacity generated by various industries, organizations, or individuals. As these demands fluctuate in capacity by time-of-day (time-zones), seasons, region, or usage patterns, public cloud providers can effectively balance and diversify all workloads to maximize utilization of the entire computing pool.

3.2. Private Clouds

Private clouds are computing resources operated, owned, and controlled by the organization that uses them internally inside the enterprise. Security and privacy, low latency, regulatory compliancy, auditing, data sovereignty, and previous investment into the IT infrastructure (such as virtualization) are among major reasons to operate in a private cloud environment and to have complete control of the IT

infrastructure. Private clouds can also include specialized computing fabrics (such as specific chips) that are not available on general purpose public clouds, and are serving the special needs of the organization.

Whether public or private, at the bottom of the cloud architecture is a physical layer that consists of physical hardware resources, such as data storage, networking, and CPUs. A virtualization layer sits above this and abstracts physical resources through virtualization technology (hypervisor) to provide virtual computing resources, such as virtual CPU, virtual memory, and virtual storage.

3.3. Cloud Infrastructure and IaaS

On top of the virtualization layer sits a cloud infrastructure layer that manages virtual resource allocations and provides fundamental capabilities for the cloud fabric, such as reliability, scalability, security, load balancing or monitoring, and metering. The infrastructure as a service (IaaS) usage model takes place here for rapid, self-service provisioning of virtual machines (VM) without a need to interact with the IT staff or a need to understand the details of the virtualization layer. An example of a public infrastructure as a service provider is Amazon's Elastic Compute Cloud (EC2). Organizations who have invested in their on-premise infrastructure can further advance toward a cloud model with either VMWare vCloud or Microsoft Hyper-V Private Cloud to enable an internal IaaS model.

3.4. Cloud Platform and PaaS

The layer above the infrastructure is a cloud platform that adds cloud capabilities to the operating system, storage, database, or other middleware or runtimes available for cloud applications. The platform as a service (PaaS) usage model can be applied here to further abstract the users from details of the base infrastructure operations (such as virtual machine provisioning) and platform maintenance (such as operating system or database updates). An example of a public platform as a service offering is Microsoft's Windows Azure platform. The Windows Azure Platform is also emerging for private cloud as "Windows Azure Appliance" to enable internal PaaS.

3.5. Cloud Application and SaaS

The top-most layer is the actual cloud application and service layer that provides business functions for the end users. It includes various SOA services, applications, content delivery services, or other cloud-related services. The software as a service (SaaS) business model is executed at this level and enables the user to focus on consumption and customization of desired business capabilities. Examples of SaaS include Bing Maps, Microsoft Office Live, or Google Docs.

4. Impact of Cloud Computing on the Geospatial Landscape

Anyone who has used Google Maps or Bing Maps has already experienced the impact of cloud computing on the geospatial landscape. Traditional GIS players and new geospatial entrants alike have enabled users to search, route, locate, geocode, and interact with spatial data through a Web browser. Easy-to-use mapping capabilities on the Internet, delivered via cloud computing architectures, are widely accepted tools for basic geospatial functionality. Proliferation of basic geospatial concepts used in online maps (such as map navigation, scale, layers, markers) enabled organizations to better utilize spatial information as people involved in processes become more skilled in working with Web maps. The concept of mashups: mixing various online data feeds and spatial content increased demand to exploit new information from existing sources by geocoding them and integrating them. Sharing information through collaborative Web applications showed value in communication of geographically diverse teams, interconnecting various communities and speeding up the information flow. With smartphone devices empowered with map applications like Google Maps or Bing Maps, users can experience another level of working with spatial content through touch screens and location based concepts. In these respects alone, cloud-based computing has already had a profound effect on the geospatial world in terms of both the awareness of geospatial and driving the industry to new capabilities and methods of delivery. In the next section, we will look at some of the most promising geospatial areas that can leverage public cloud computing technology.

4.1. Publishing Spatial Data on the Cloud

Spatial data are large, storage intensive, and require specialized types of organization and optimization (or caching) to quickly retrieve required data at the given detail, area, and layer. Moving or copying these datasets from producer to consumer is time intensive and may require specialized administrative techniques, such as incremental updates. The ability to access spatial information from the source offers a more efficient process and avoids the local copying of data.

Spatial data infrastructure (SDI) has long been put forward as a concept, architecture, and set of standards for how spatial data can be remotely searched and accessed. The geospatial sector, through organizations such as the Open Geospatial Consortium (OGC®) and ISO, has defined standards for interoperable spatial data and architecture based on SOA principles. Widespread adoption of OGC standards and various SDI initiatives have prepared the geospatial industry to embrace cloud computing technologies for hosting spatial data and spatial processing services. Going forward, the combination of public clouds and SDI concepts could provide the best architecture to realize the promise of SDI. With cloud computing, SDI can reach the necessary scalability and conformity to meet performance rules of SDI frameworks, such as INSPIRE. For example, the INSPIRE directive specifies that the view service must download images of 470 KB size within five seconds and must handle up to 20 simultaneous requests per second. Further, INSPIRE services must have an overall 99 percent availability.

4.2. Processing Spatial Data in the Cloud/High-Performance Computing

Many geospatial processes are intensive and thirsty for computing resources. This ranges from geocoding, routing, and reprojection to analytical applications, such as network tracing, geospatial analysis, spatial statistics, pattern recognition or rendering, and visualization tasks. While the effectiveness of spatial data processing depends on many factors, such as workload patterns, data volumes, partitioning, and I/O performance between storage and processing units, the massive elastic scalability and parallel processing foundation of the public cloud offers real promise for high-performance computing (HPC) in general and large-scale geospatial processing in particular.

When coupled with a “pay as you go” business model, public cloud computing offers the potential for organizations to perform geospatially based analyses previously beyond the scope of their processing capabilities. A challenge for geospatial software vendors is to make the necessary enhancements to their architectures to take advantage of these emerging HPC environments in the cloud.

4.3. Backend Services for Mobile Applications

Cloud computing is uniquely positioned to support mobile devices. Through centralized processing in a public or private cloud, a smart device with a relatively small on-board application can take full advantage of the power of the processing and data available in the cloud. According to a recent Morgan Stanley report, mobile access to the Internet using handheld devices is ramping nearly 10 times faster than the equivalent desktop Internet ramp. Further, Morgan Stanley predicts the number of mobile users accessing the Internet will surpass the number of desktop users within five years. Google Maps on the Blackberry or iPhone demonstrate the utility of these types of applications. In the future, we’ll see the arrival of many innovative mobile geospatial applications leveraging the processing power of the cloud connected to an intelligent mobile device with integrated positioning technology, such as GPS. Location-aware applications in areas such as intelligence-led policing and augmented reality will become mainstream.

4.4. Geospatial Complex Event Processing in the Cloud

Capturing location-enabled event data from the Web, sensors, mobile devices, or other sources creates a need for a robust infrastructure that can handle and sustain occasional peak loads, such as during a time of emergency or crisis. Cloud computing can offer the needed elasticity to process, store, and further analyze discrete events. Such complex processing may include filtering and aggregation of events, pattern or trend detection, and evaluation of spatial rules for location-aware events. The resulting complex event may be further weighted for its relevance and sent to the right actuator, application, or service for subsequent processing. Integrating incoming complex events with other external spatial data gives organizations actionable intelligence and visualized results for optimal decisions in near-real time. Cloud computing ensures sustained performance for event processing during fluctuations in demand.

4.5. Inhibitors to Cloud-based Solutions

Many of the inhibitors today to public cloud-based implementations are not technical. They typically revolve around concerns about custodial issues of data security, and in some cases, specific regulations that make it difficult to allow data to be held or processed outside the physical boundaries of an organization or the boundaries of a given geography. Cloud-based providers are responding to the geographical concern and have provided mechanisms for organizations to stipulate where their data are physically stored in the cloud (e.g., in Europe or North America). As for policies which do not allow data outside of the physical domain of an organization, some of these constraints will ease as cloud computing grows in acceptance. However, some industries and sectors will require a private cloud approach.

5. Intergraph's Cloud Computing Strategy

Intergraph fully anticipates offering cloud-based, hosted solutions in markets and areas that make sense for us and our customers. We are taking incremental steps and making technology investments that will enable us to fully leverage cloud models in the future. We look on cloud computing from a holistic perspective. This includes not only hosting options, but other areas, such as application architecture, integration of cloud services with current front-ends, new business and usage models, optimizing management and licensing models, flexibility of re-hosting between private and public cloud, and mobile device or sensors capabilities and their access to cloud computing.

Intergraph divides its cloud computing strategy into the following areas:

5.1. Application Architecture for Cloud Computing

Cloud computing imposes a need for service oriented architecture. SOA is an architectural style for distributed computing based on the principles of loosely coupled, stateless, and self contained business capabilities. During the past several years, Intergraph has invested into evolving our application architecture to SOA principles. We continue on our overarching plan and vision of the service-centric approach based on four primary tenets: Integration, hosting, technology, and discovery (Figure 3). With this concept, cloud computing presents an efficient hosting technology that Intergraph services can leverage. With cloud computing's highly elastic scalable environments, we understand a need to focus on applying the right design patterns for parallel execution, dynamic scaling, performance metrics, and security of our services.

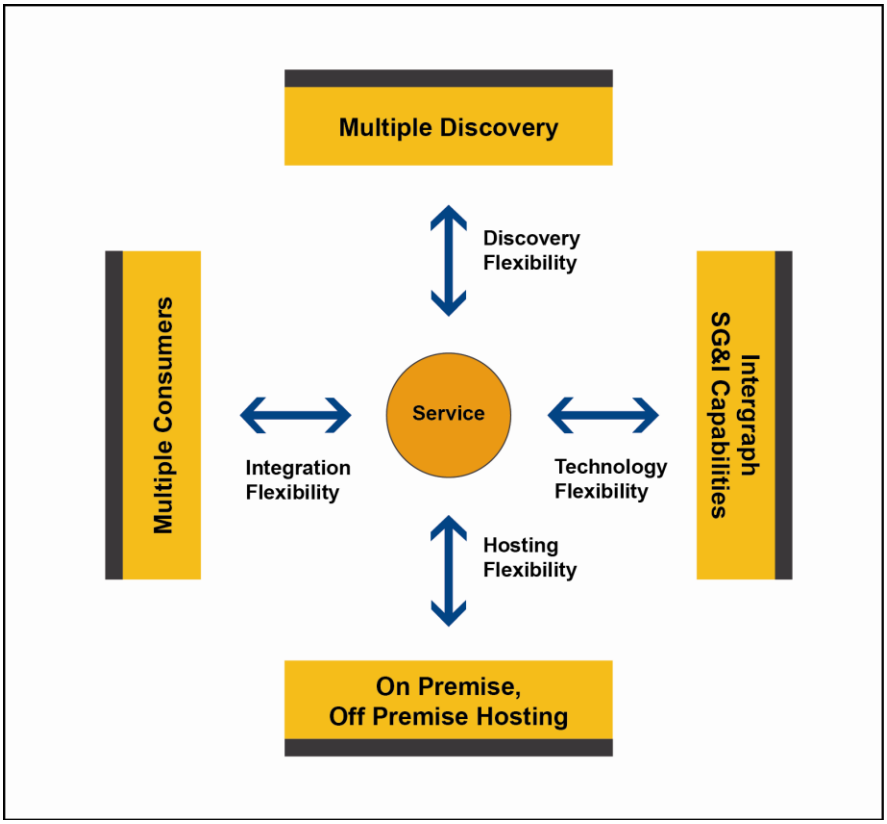

Figure 3: Intergraph's service-centric approach is based on four tenets of flexibilities: Integration, hosting, technology, and discovery.

We are continuing to make a substantial investment in next-generation geospatial technology based on SOA principles and employing the latest Microsoft .NET ecosystem. These investments will allow us to fully exploit the Microsoft Azure platform for geospatial application development. As Azure has evolved over the past years, it has grown to not only include a public cloud-based approach, but also a private cloud-based component with the upcoming Microsoft Azure Appliance. This dual public and private cloud model makes a lot of sense for customers constrained to a private cloud approach, or who want to start with a private cloud and transition to a public cloud over time.

5.2. Certifying on Existing Cloud Platforms

Pricing, reliability, and many other attractive features of today's public cloud offerings are causing us and our customers to examine how we can leverage the cloud to provide geospatial-specific platforms for our customers to augment what they are doing with their own computing resources. We are in the process of certifying GeoMedia® WebMap for the Amazon EC2 environment. This will allow customers to implement GeoMedia WebMap and related vertical business applications in a fully-certified fashion in the Amazon environment. We are also in process of evaluating Microsoft Azure and SQL Azure for its ability to host GeoMedia WebMap and spatial data in a public cloud environment. Meanwhile, Intergraph will continue evaluating other existing products for certification in cloud-based environments. Products such as GeoMedia SDI Professional and GeoMedia SDI Portal could be early candidates for certification on either or both Amazon EC2 and Microsoft Azure. Through a pilot project in Denmark, we will offer other Intergraph applications under the SaaS model based on the Amazon EC2 infrastructure.

5.3. Cloud Management and Licensing Tools for Geospatial Solutions

Intergraph understands the need for self-service front-ends, for both the management and provisioning of new computing resources in the cloud environment. In this respect, we are focusing on providing better administration tools to manage multiple instances of our geospatial platform, applications, or services and straightforward methods to manage licenses or charges per use. This is in line with cloud monitoring tools that helps end-users better understand the usage vs. pricing model.

5.4. Cloud Services Consumption and Integration

Visualization and consumption of the results coming from cloud computing is a key capability of our desktop and Web solutions. For example, several Intergraph products already incorporate access to Google Maps and Bing Maps within their application framework, as well as OGC data sources, such as WMS and WFS. However, consuming data is just a first step. The next challenge is to provide easier access to the existing processes and analysis available in the cloud, and moreover to provide a convenient way for the end user to define new ad hoc custom processes and process flows, while working in a familiar productive environment. As an example of this vision, we have successfully demonstrated a prototype that allows the end user in GeoMedia to consume standard OGC WPS processes, visualize them, define new processes, and run them in the cloud – all visually chained using Microsoft Workflow Foundation integrated into GeoMedia as Figure 4 illustrates.

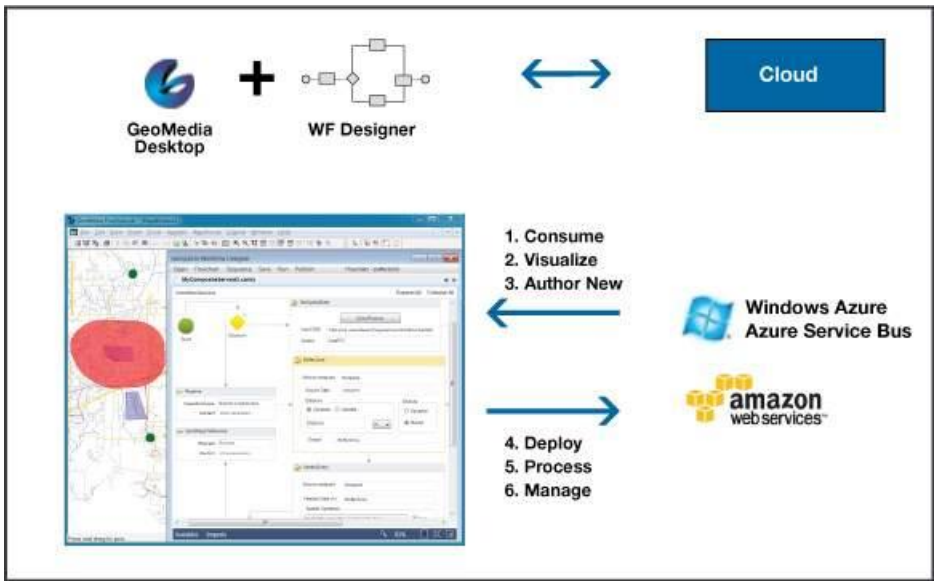

Figure 4: This diagram illustrates a prototype of Integrated Workflow Designer inside GeoMedia, which allows user-friendly access to existing services in the cloud, author visually new, and run them in the public or private cloud.

Through this type of integration, Intergraph customers will be able to take advantage of existing cloud-based processing and data sources within their familiar environment.

In addition, Intergraph has also invested into its own hosting infrastructure technologies, such as Terramapserver or Respublica Intranet, which many European customers use today for running their off-premise geospatial applications, services, or data with high availability and security. Both dynamic data centers, located in Munich and Vienna, offer various business models, ranging from server housing to the full software as a service (SaaS) model. Today, we are harvesting experiences gained by operating these data centers, understanding customer needs, business models, and legal constraints to consider moving them further into the realm of the cloud computing.

6. Conclusion

Cloud computing is a growing trend, based on sound computing principles for distributed computing. Its influence on application architecture and IT deployment strategies will continue to accelerate for many reasons, both economic and technical. Public clouds offer reasonably priced, elastic computing resources with truly on-demand consumption models that users can immediately leverage without capital expenditure in IT infrastructure. Capabilities delivered through cloud computing, such as smart server utilization, automated management of IT operations, on-demand elastic scalability, and SaaS, will become an important part of IT strategies for many organizations

Data security and compliance capabilities are improving in public clouds providing viable options even for organizations restricted by regulatory policy. At the same time, legislation constraints will modernize over time to reflect new technology options, which has happened in the past with the Internet.

At Intergraph we closely watch and are researching new possibilities of cloud computing from its various technical aspects, applicability to existing products, and business value proposition. We see infrastructure as a service as a transitional step toward platform as a service that further abstracts core infrastructure by providing and managing the foundation software platform and continuing to reduce operational cost.. Microsoft Windows Azure is a promising technology rapidly evolving to deliver the full promise of the cloud computing benefits not only in the public, but also for private clouds.

Intergraph's strategy to leverage the cloud in its products includes four major areas:

1. Cloud computing application architecture based on SOA
2. Certifying products on existing cloud platforms
3. Providing cloud management tools, including cloud friendly licensing policy
4. Providing cloud services consumption and integration within existing products

Intergraph will continue to take steps that will allow our customers to take advantage of the computing platform of their choice based on business demand, technology, architecture, and standard readiness.

For more information about Intergraph, visit our website at www.intergraph.com.

Intergraph, the Intergraph logo, and GeoMedia are registered trademarks of Intergraph Corporation. Microsoft and Windows are registered trademarks of Microsoft Corporation. Other brands and product names are trademarks of their respective owners. Intergraph believes that the information in this publication is accurate as of its publication date. Such information is subject to change without notice. Intergraph is not responsible for inadvertent errors. ©2011 Intergraph Corporation. All Rights Reserved. 1/11 SGI-US-0035A-ENG