

THE INCREASING IMPORTANCE OF THE PUBLIC CLOUD

Not long ago, cloud computing was considered a disruptive technology, embraced primarily by innovative startups and more audacious, boundary-pushing businesses. How quickly that has shifted: Today, enterprises in even the most traditional of industries are migrating to the cloud, some of them completely leaving their legacy systems behind.

In the world of database management systems (DBMS), Gartner numbers show that cloud services represented [\\$10.4 billion of the \\$46.1 billion DBMS market in 2018](#), not including hosting-related licenses. Gartner also notes that the overall DBMS market grew at 18.4% from 2017 to 2018—its highest growth rate in more than a decade. Cloud-based DBMS accounted for an astounding 68 percent of that growth.

The numbers are huge, and yet they shouldn't be a surprise. Data drives the digital economy, and every organization has an increasing need to easily store, access, analyze, share, and secure their digital assets. The cloud offers storage capacity, availability, scalability, and agility on a scale that would be resource- and cost-prohibitive to maintain in an on-premise environment, while it dramatically reduces the overall complexity of data management.



No matter the application, the cloud offers some undeniable advantages over on-premise IT infrastructure:



**COST
SAVINGS**

One of the strongest motivators for enterprises to move business-critical applications to the cloud is simple finances. Because the cloud service provider takes on the cost of procuring and maintaining hardware and other infrastructure, the enterprise is able to reduce or eliminate hardware purchase costs, as well as the associated costs of space and power to house and operate servers. Adopting cloud solutions also reduces labor costs, as the management responsibility is passed on to the provider and the business can reallocate or reduce their in-house IT staff. Additionally, cloud solutions typically operate on a pay-as-you-go pricing model. Whether you're purchasing server capacity or user licenses, you pay only for what you need. This eliminates or reduces expenses associated with overprovisioning on-premise resources.



**EASIER
SCALABILITY**

Cloud solutions allow businesses to easily and rapidly scale storage and computing resources according to their business requirements. They can provision as much as they need with the click of a button, enabling them to offer additional services and products more quickly and temporarily scale up resources during peak seasons to accommodate greater demand.



**GREATER
AGILITY**

As mentioned above, the cloud's improved scalability allows for the rapid provisioning of computer resources, compared with the weeks (or months, in some cases) it takes to procure and provision on-premise resources. The pay-as-you-grow model allows new technologies to be easily trialed. This greater agility allows the company to move more quickly to introduce new services and new applications into the marketplace.

HYBRID CLOUD: HOW ON-PREMISE AND CLOUD TECHNOLOGIES CAN WORK TOGETHER

Moving workloads to the cloud has helped many enterprises lower operating costs while simplifying and modernizing their architecture. The fluidity of cloud environments enables more flexibility and agility across business operations, greater availability of resources, and faster time-to-market.

But for many other companies, significant technological, operational, financial, and security hurdles stand in the way of fully moving their assets to the public cloud. With millions of dollars already sunk into huge legacy systems, many enterprises are understandably reluctant to completely divest themselves of their longstanding infrastructure to go all-in on the public cloud. Many companies also have an aversion to moving mission-critical systems to the cloud because of lingering concerns over privacy and security. For industries bound by compliance regulations, such as healthcare and finance, transmitting data over a network that could potentially be accessed by third parties is sometimes considered too great a risk.

At the same time, these organizations typically recognize the advantages of the cloud, and historically reluctant businesses such as major banks, telecommunications firms, and other large organizations have been making incremental moves to the cloud by employing a hybrid cloud model. The term “hybrid cloud” refers to the use of one or more public cloud providers alongside internal resources—either a private cloud system or the company’s on-premise IT infrastructure. This configuration allows businesses to store and secure sensitive data internally while

leveraging the efficiencies of the public cloud as needed. (Contrast this with “multicloud” systems, which is simply the adoption and simultaneous use of more than one cloud computing system.)

Hybrid cloud technology offers several benefits to these types of users, particularly because it gives IT more control over how and where data is stored than in a straight public cloud platform. Using a strategy called “cloud bursting”, companies can run typical workloads on-premise while purchasing extra computational resources through the public cloud during times of heightened demand, such as tax season or end-of-year reporting. Companies can also use hybrid cloud to develop and offer new services in the cloud, while slowly migrating older on-premise services to the cloud over time, thus reducing the likelihood of service disruptions.

Banks, for example, still process about 90% of credit card payments via on-premise mainframes. Regulations also prevent them from storing sensitive data in the cloud because of the risk of exposing confidential customer information to third parties. A hybrid cloud approach allows them to maintain compliance with regulatory frameworks by keeping CRM and billing systems within their private on-premise infrastructure, while still allowing them to leverage the cloud for innovation and fast delivery of new products and services.



Globalization is also driving many industries toward a hybrid cloud model. General Data Protection Regulation (GDPR) data location compliance requires companies that collect data in the EU to store it in the EU. However, a U.S.-based multinational with a data center located in the U.S. may not be willing or able to build a separate data center overseas to comply with this regulation. In this scenario, the company can store its EU customer information on cloud servers located in the EU to satisfy GDPR requirements.

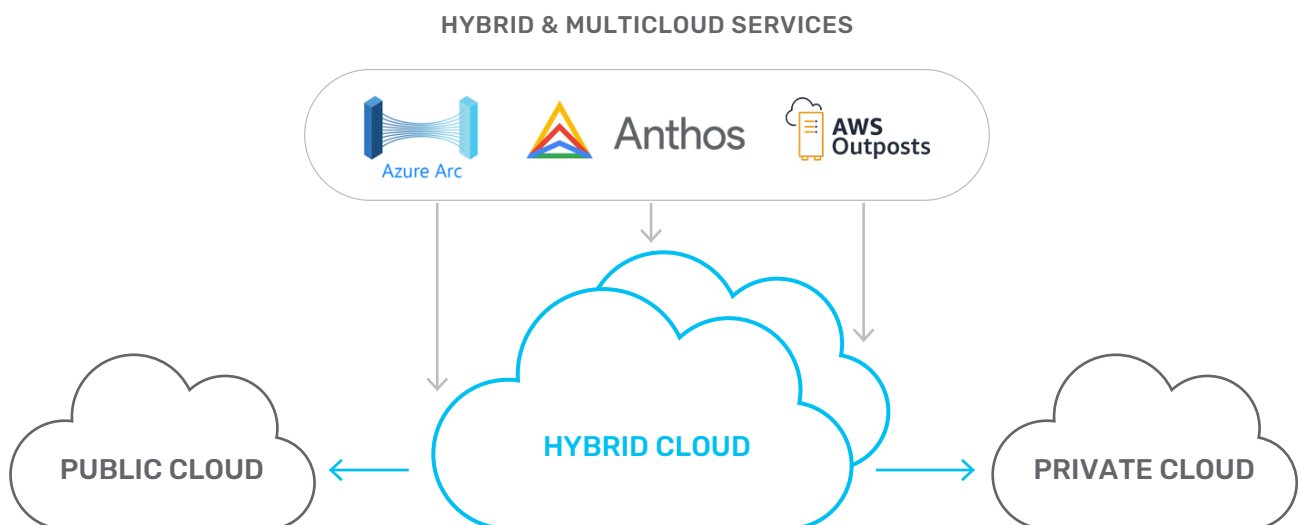
In the last year, we've seen the three top public cloud providers—Google, Microsoft, and Amazon—respond to this trend by introducing hybrid or multicloud services. The common goal of all three is to provide enterprises with a frictionless transition to the public cloud by delivering a consistent experience across their on-premise and cloud environments.

Google's Anthos is significant for being the tech giant's official entry into the enterprise data center and the first official multicloud platform from a mainstream public cloud provider. Anthos is powered by one of the most respected open-source platforms, Kubernetes, via the Google Kubernetes Engine (GKE), the managed containers-as-a-service offering on the Google Cloud platform.

A few months after Anthos was unveiled, Microsoft announced its own hybrid and multicloud platform, Azure Arc. This new platform allows users to deploy Azure services anywhere and extends Azure management to any infrastructure.

While Microsoft and Google have embraced hybrid cloud and multicloud strategies, Amazon has historically encouraged customers to fully adopt the public cloud. Amazon signaled a strategic messaging shift when it announced AWS Outposts, a platform that allows customers to take advantage of AWS's popular storage and analytics features both in AWS Cloud and within customers' private legacy data centers. Customers can choose from two variants: Native AWS and VMware Cloud.

Another indicator of the hybrid and multicloud trend is the development (and IBM acquisition) of Red Hat OpenShift. This enterprise-ready Kubernetes container platform provides the necessary tools for managing hybrid and multicloud deployments on both public and private cloud systems.



NOT EVERY CLOUD IS THE SAME: THE BENEFITS OF A MULTICLOUD STRATEGY

Even though hybrid cloud is on the rise, an organization rarely finds that “one cloud fits all.” More often, companies are compelled to adopt a multicloud strategy, leveraging multiple cloud services. There are several drivers behind this.

The most common reason organizations partner with multiple cloud providers is that no single cloud platform offers everything they need. A close look at some of the major cloud solutions reveals some stark differences in features and capabilities. Both Amazon AWS and Microsoft Azure, for example, include frameworks for building, training, and deploying machine learning (ML) models—SageMaker for AWS, and Machine Learning Studio for Azure. But they each take different approaches to accomplishing their goals. Machine Learning Studio uses a visual drag-and-drop interface, while SageMaker relies heavily on code and requires a deeper knowledge of data engineering.

As noted above, multinational organizations have the additional consideration of location-based regulations when considering cloud technologies. A company opening a factory in China, for example, will be required by law to store its data with a Chinese cloud service provider. By the same token, an American company storing its data on Google servers in Ireland is still subject to subpoenas by the U.S. government under the [CLOUD Act](#) because Google is an American company. Each cloud service provider has a different geographical spread that spans dozens of separate regions, which further complicates matters.

Of course, in any situation, vendors may discover that their contracted cloud provider simply can't accommodate their myriad business requirements. Yet, they'll remain “locked in” due to the high transfer costs and technical challenges of moving to another provider.

Most organizations eventually choose to work with multiple cloud providers. In a recent survey of public cloud users, 81% of respondents said they are working with two or more providers.



Multicloud isn't a bad thing, mind you. In fact, there are several benefits to taking a multicloud approach:



MORE FEATURES

By placing workloads across multiple cloud environments, enterprises can take advantage of different feature sets and tailor services to meet varying business needs, letting them innovate on their terms. This allows them to avoid vendor lock-in in many cases.



MORE LOCATION OPTIONS

Using multiple cloud providers can expand an organization's regional reach, enabling them to position data centers closer to their users, which can improve performance and reduce latency. It also helps solve data locality challenges created by geographically-based regulatory requirements as previously discussed; having the cloud server in the same country as the user makes it easier to comply with local laws and regulations.



REDUNDANCY

All cloud providers experience [downtime issues](#), even if these instances are minor. A multicloud setup ensures high availability and resilience when data is available on different clouds.



COST EFFICIENCY

Working with multiple cloud providers enables companies to take advantage of the best available pricing. For example, storage at Azure Blob costs less than it does at AWS S3. This can protect companies against vendor-specific price hikes and allow them to negotiate better deals from each vendor over time.

That said, while a multicloud approach offers enterprises significantly more flexibility, it also injects greater complexity into the computing environment. In multi-region and multicloud deployments, the replication and sharing of data among multiple, geographically-distributed active clusters to support global business operations is particularly challenging. **In many cases, an enterprise can benefit from having a data fabric that connects various clouds and leverages the benefits of each of them.**

WHAT CHALLENGES DO HYBRID AND MULTICLOUD DEPLOYMENTS FACE?

Hybrid and multicloud approaches bring some much-needed flexibility to companies' data management strategies. However, they also introduce new challenges and can create inefficiencies that need to be addressed.

Organizations may face some or all of the following concerns:



LACK OF DATA LOCALITY

Time- and data-intensive applications and services require data locality, in which data is stored as close to the computation as possible. This can be challenging, especially when replicating data across multiple environments and cloud vendors while ensuring consistency and high concurrency.



DATA REPLICATION OVERHEAD AND NETWORK COSTS

Most cloud providers include the cost of data ingress as part of the service itself but charge a separate data transfer fee to move data out of the cloud. Costs range between 5 and 20 cents per GB every time data is moved and can contribute up to [30% of total fees](#). With enterprises moving terabytes of data every day that can add up to an unexpectedly monstrous bill.



DATA PRIVACY AND SECURITY ISSUES

In some cases, data may not be suitable for storage on the cloud at all. While hybrid cloud deployments typically let you mask and filter sensitive data and store it off of the public cloud while still letting you leverage the cloud for "low risk" operations, different regional regulations may require the implementation of separate workflows in various areas.



SERVICE LEVELS AND AVAILABILITY CONCERNS

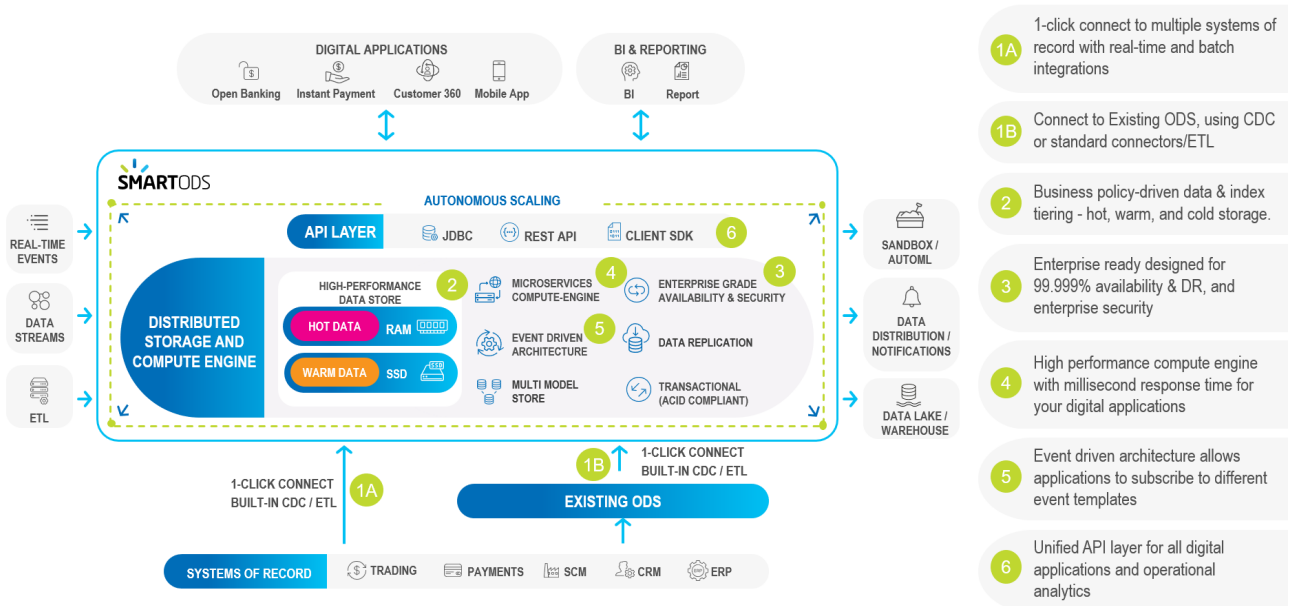
Ensuring service levels and high availability is a challenge that is not dependent on the deployment environment, however a multi-cloud environment can complicate this. At peak periods with heavy data transfer volume or a large number of concurrent users—such as Black Friday, Covid-19 pandemic surges, and the U.S. Presidential elections—response time can be degraded and servers can crash altogether. Cloud environments need to be able to automatically scale up and out to meet CPU and RAM requirements in order to avoid these issues and ensure customers are not dissatisfied.

GIGASPACE'S INSIGHTEDGE: BUILDING A SUCCESSFUL HYBRID AND MULTICLOUD STRATEGY

To address the challenges discussed above, GigaSpaces offers a suite of modern data and analytics platforms.

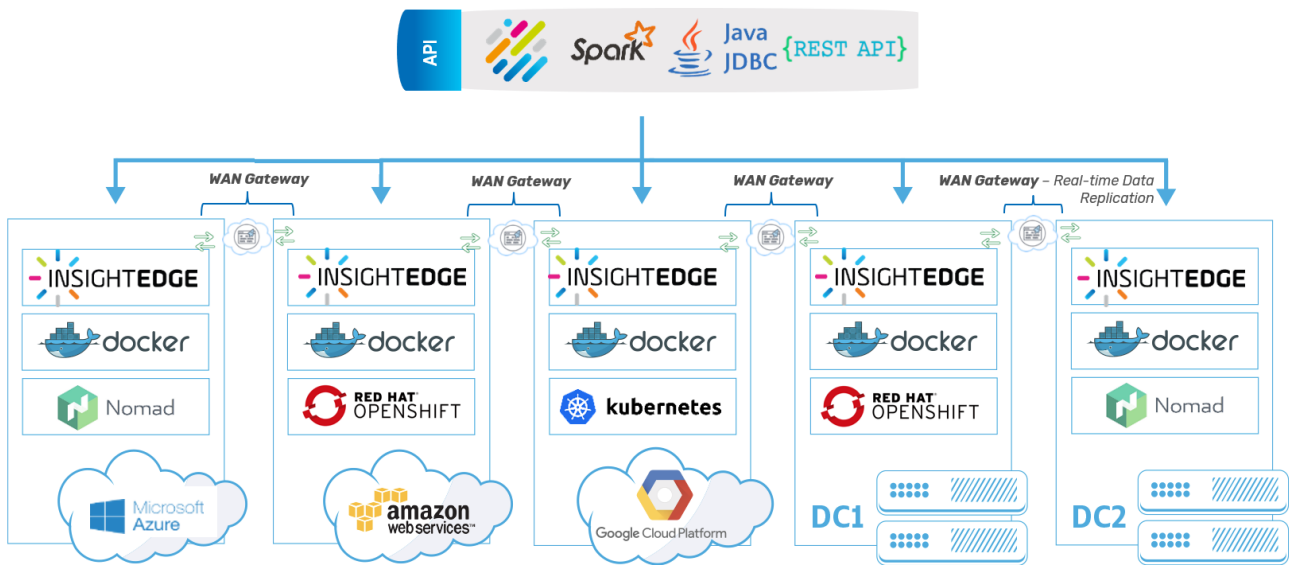
InsightEdge was designed to help enterprises seamlessly accelerate and scale their mission critical, time-sensitive applications and services. This distributed in-memory software platform can ingest, process and store large volumes of any data type, and ensures low latency performance dynamic scale across all environments. Customers leverage InsightEdge to power open banking initiatives, real time fraud and risk analysis, customer 360, analytics and BI on fresh data and more.

The InsightEdge platform integrates to siloed operational data stores and to the enterprise's systems of record with a single click (no-code connect), aggregating data in a low-latency data fabric. The dynamically scaling [Data Integration Hub](#) offloads API access from the enterprise's data stores (on-premise and cloud) and delivers rapid application response times to internal and external customers no matter the load while ensuring always-on services. The ability to colocate business logic with data in memory and perform dynamic server-side aggregations, reduces the movement of data to the client, and delivers extreme performance.



Its out-of-the-box data replication efficiently moves data between on-premise data centers and the cloud, among different cloud vendors, or among multiple regions within a single cloud platform. In hybrid cloud deployments, InsightEdge's Hybrid Cluster functionality supports a unified cluster between on-premise and the cloud, which eliminates the need for data movement.

InsightEdge is designed to minimize the amount of data that is transferred during replication by transferring only the changes made to the data rather than the entire data store. Data can also be filtered for privacy purposes.



The advanced data management, processing, and storage capabilities of InsightEdge offer several distinct advantages.



NETWORK EFFICIENCY

Because InsightEdge allows organizations to replicate only changes to data or aggregate volumes of data, transfers are faster, tie up less bandwidth, and incur lower data transfer fees. Data transfers can also be reduced via the application of various compression methods.



DATA LOCALITY

InsightEdge allows companies to replicate data from one region to another in seconds or minutes. Policies can be configured to make replication between data centers unidirectional or bidirectional. Regardless of how the data is distributed, the platform unifies the data and allows it to be analyzed as if it were generated locally. External databases can be connected to InsightEdge with a single click, without having to write any code.



PRIVACY AND COMPLIANCE

InsightEdge allows organizations to define what data fields should be anonymized or encrypted during storage and replication. Policy configurability is based on a P0/P1/P2/P3 scale: P0 data should stay on-premise, P1 should be encrypted when replicated to the cloud, P2 should be anonymized, and P3 may be stored in clear text. Each customer can bring their own encryption keys and key size, allowing them to control the level of encryption.



AI-DRIVEN AUTONOMOUS SCALING

InsightEdge allows businesses to size the data fabric of their infrastructure for typical business demands while enabling them to automatically handle peaks. The platform can detect spikes based on CPU and RAM utilization and automatically increase memory, CPU, storage, and other resources for each service and subset of the data experiencing heightened demand, with no impact on performance or loss of service. Similarly, GigaSpaces can automatically scale resources down to save infrastructure costs during off-peak hours or when demand is low.

Scaling alert settings [X]

> CPU Factor 4 Scale up at 80% Scale down at 10%

▼ RAM

Scaling factor Custom

2 4 8 16 32 3

Scale up Trigger at Time Period
× Factor 80 % 00:02:00

Scale down Trigger at Time Period
÷ Factor 10 % 00:02:00

Cancel Apply



RESILIENCY AND AVAILABILITY

InsightEdge supports both active-active and active-passive deployments for optimal availability and resilience and supports auto-recovery in the event of a failure. This redundancy allows companies to keep their workloads on-premise with a disaster recovery site in the cloud or have clusters deployed in data centers in two different geographic locations (or zones within the same region), directing users to each based on proximity.



CLOUD BURSTING

InsightEdge can expand resource capacity through “cloud bursting.” Businesses can run workloads in their on-premise infrastructure and leverage the elasticity of the cloud when they need to run more intensive calculational business processes, such as end-of-period reporting. Bursting to the cloud for additional capacity for finite periods reduces strain on on-premise systems and lowers the expense of provisioning additional hardware for limited usage.

WITH INSIGHTEDGE ENTERPRISES CAN ACHIEVE



AGILE

introduction of new digital applications



OPERATIONAL

real-time reporting, analytics and BI on fresh data



MULTI-REGION

data locality across data centers and cloud regions



MODERNIZATION

and offloading legacy systems (Mainframe, AS/400)



ACCELERATION

of applications and batch processes for online services and time-sensitive data



CLOUD NATIVENESS

to develop once, and deploy on any environment – hybrid and cloud

TYPICAL INSIGHTEDGE USE CASES

INSTANT PAYMENTS & FRAUD PREVENTION

The shift to cashless commerce and real-time mobile payments is transforming the way we pay, the services we use and our expectations as customers. To ensure a positive customer service experience, service providers must be able to process payment transactions in real-time, prevent fraud and ensure regulatory compliance.

Accurate and fast prevention of fraudulent payments, trading and deposits require low-latency response time to complex queries even during unplanned peaks and at high concurrency. Event-driven analytics triggers and notifications to applications trigger on-the spot analysis and workflows against fraud as it is happening.

CONTINUOUS RISK ANALYSIS

Reducing risk exposure and meeting regulatory requirements, while lowering operational costs, are still key challenges to financial services and insurance organizations. By ingesting, processing and analyzing billions of messages in real-time into a live risk result store, hundreds of applications can concurrently leverage machine learning models at sub-second speeds while maintaining transactional integrity and enabling continuous risk analysis.

CUSTOMER 360 & HYPER- PERSONALIZATION

Intentions and past activities are facilitated by ingesting real-time customer interaction data at high throughput and analyzing the data at the moment it is generated. Enriching these insights with historical data, combined with sub-second response times, powers highly-personalized and time-sensitive customer services. Typical examples include personalized offers, location-based marketing, opening of new accounts, purchasing of insurance policies, loan requests and more. In short, making customers happy. Furthermore, deep learning and predictive analytics help in the effective identification of the likelihood of churn and management of upselling.

DYNAMIC PRICING

To be price competitive, all online offers must react quickly to competitor promotions and other influencing parameters, depending on the industry. For example, the transportation business entails many complex, interdependent parameters such as order requests, available drivers/pilots, customer geographical location, weather, holiday peaks and more. Retail and eCommerce can leverage machine and deep learning models for predicting pricing trends, generating on-the-spot offers and changing consumer considerations about what they are purchasing, in a moment. To achieve this, data from multiple sources must be ingested, processed and analyzed in real-time-leveraging the current data as well as historical information - for the smartest insights and ability to proactively ensure the best efforts.

OPEN BANKING INITIATIVES

Existing traditional and disparate architecture cannot adequately and efficiently support the exponential increase in queries that are required for open banking. You can offload the systems of records and operational data stores with an always-on, low latency and dynamically scaling data fabric that provides a unified API layer for agility to develop new applications to monetize on the Open Banking APIs.

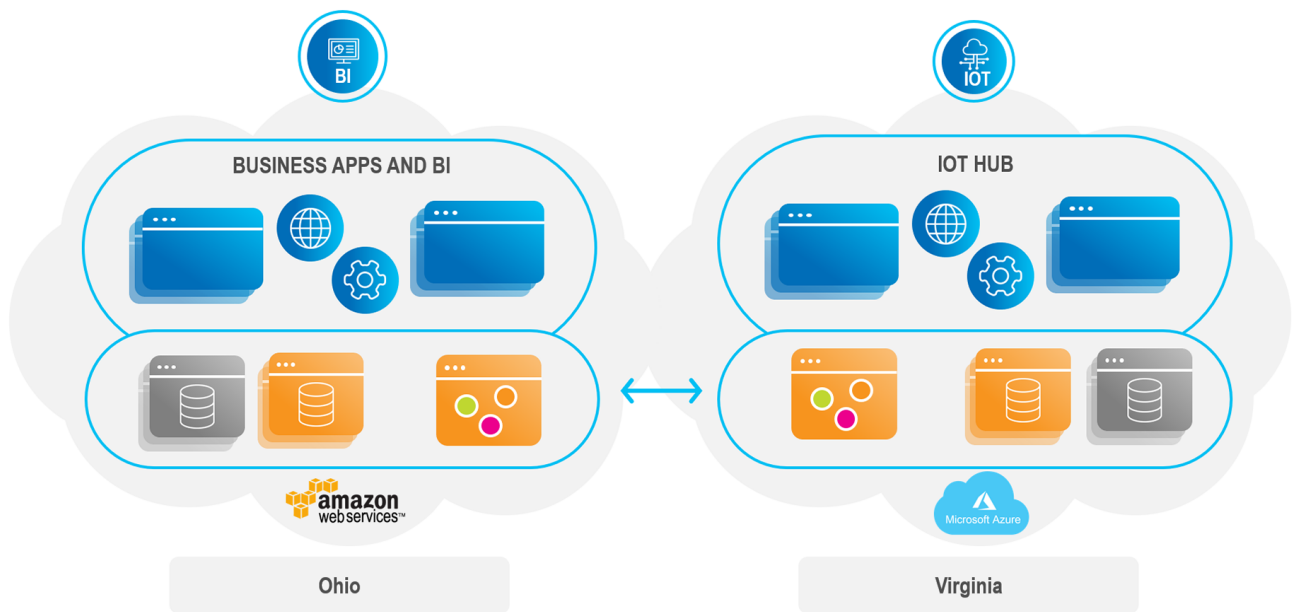
OPERATIONAL BI ON FRESH DATA & BI ACCELERATION

Live-fresh data can be queried across any dimension in sub-seconds and at high concurrency (100s of analysts/apps). Long batch processes can be accelerated from hours to minutes.



CASE STUDY: GIGASPACES IN THE MANUFACTURING INDUSTRY

Modern enterprises require the ability to replicate their mission-critical data across applications, regions, clouds, and data centers. In this case, a leading manufacturing company needed to ensure real-time data replication of its IoT data from its manufacturing plants in Virginia to its data analysis and BI divisions in Ohio for analysis and of the data that is leveraged for predictive maintenance. One of the key components of the InsightEdge solution was the ability for the customer to work both synchronously and asynchronously between its Azure IoT Hub (in Virginia) and AWS (in Ohio).



The GigaSpaces solution combines synchronous and asynchronous product features, and by doing so, supports hot and cold data accessibility for multi-region and multicloud deployments. Here's how the architecture works.



The following 3 steps were implemented:

1. Set up 1...n sites globally

pu.xml WAN Gateway Azure side configuration snippet (Sink):

```
<beans
.....
<os-gateway:sink id="sink" local-gateway-name="Azure-us-east-2" gateway-
lookups="gatewayLookups" local-space-url="jini://*/*/SpaceAzure-us-east-2" start-embedded-
lus="false">
  <os-gateway:sources>
    <os-gateway:source name="AWS-us-east-2"/>
  </os-gateway:sources>
</os-gateway:sink>
.....
</beans>
```

pu.xml WAN Gateway AWS-us-east-2 side snippet (Delegator):

```
<beans
.....
<os-gateway:delegator id="delegator" local-gateway-name="AWS-us-east-2" gateway-
lookups="gatewayLookups" start-embedded-lus="false">
  <os-gateway:delegation target="Azure-us-east-2" />
</os-gateway:delegator>
.....
</beans>
```

2. Set up communication between the sites using the GigaSpaces LRMI protocol

pu.xml web API Azure-us-east-2 side configuration snippet:

```
<beans
.....
// defining LRMI connection to local Azure-us-east-2space
<os-core:space id="Azure-us-east-2IJSspace" url="jini://*/*/SpaceAzure-us-east-2" />
<os-core:giga-space id="Azure-us-east-2Space" space="Azure-us-east-2IJSspace" />

// defining LRMI connection to remote AWS-us-east-2 space
<os-core:space id="AWS-us-east-2IJSspace" url="jini://*/*/SpaceAWS-us-east-2" />
<os-core:giga-space id="AWS-us-east-2Space" space="AWS-us-east-2IJSspace" />

<!-- AWS-us-east-2 space based remoting services -->
<bean id="iotDataRemotingService"
class="org.openspaces.remoting.ExecutorSpaceRemotingProxyFactoryBean">
  <property name="gigaSpace" ref="AWS-us-east-2Space" />
  <property name="serviceInterface"
value="com.gigaspace.fdal.service.remoting.IIoTDataRemotingService" />
</bean>
.....
</beans>
```

3. Based on the LRMI communication channels, we can now transfer the requested data according to its priority:

High priority (Hot Data) – Get the data from the local space. If it does not exist, proceed to get it synchronously from a site which has the needed data.

RegionController java class implementation code snippet, taken from Azure-us-east-2 web API:

```
Region[] regions = Azure-us-east-2Space.readMultiple(new Region());

if (regions != null && regions.length > 0) {
    LOGGER.info(regions.length + " region objects have been read from the Azure-us-east-2space.
    Notifying IoT...");

    IoTAsyncGetRequest iotNotifyRequest = new IoTAsyncGetRequest(url, headers, user.getId(),
    Region.class);

    usSpace.write(iotNotifyRequest);

    return ok(Arrays.stream(countries).map(Region :: getProperties).toArray());
}
else
{
    Response<Region[]> regionsResponse = iotDataRemotingService.load(new Request(url, headers),
    Region.class);
    LOGGER.warning(String.format("%d region objects have been read from IoT. IoT request: GET %s",
    regionsResponse.getEntity().length, url));
    return regionsResponse.toRestResponse();
}
```

Low Priority(Cold Data) – Get the data from a remote site asynchronously according to the user profile predicting that this data will be needed for him later on.

IoTDataAsyncService java class implementation code snippet, taken from AWS-us-east-2 space business logic:

```
@EventDriven
@Polling
public class IoTDataAsyncService {
    .....
    @EventTemplate
    public SQLQuery<IoTAsyncGetRequest> template()
    {
        SQLQuery<IoTAsyncGetRequest> query = new SQLQuery<>(IoTAsyncGetRequest.class, "");
        query.setRouting(routing);
        return query;
    }

    @SpaceDataEvent
    public void eventProcess(IoTAsyncGetRequest request) {
        .....
        if (Boolean.TRUE.equals(request.getSaveToSpace())) {
            Response<?> response = iotDataService.load(url, requestHeaders, request.getEntityType());
            if (request.isSessionData()) {
                PrivateData[] data = (PrivateData[]) response.getEntity();
                sessionDataManager.write(AWS-us-east-2Space, data, request.isReplicable());
            }
            HttpStatus = response.getStatusCode();
        }
    }
}
```

By getting the user details (profile) on the first access, we can now collect all of the user's predictive data asynchronously and this data will wait for him locally once he requires it.

ORCHESTRATION BENEFITS WITH GIGASPACEs AND IBM RED HAT OPENSIFT

GigaSpaces InsightEdge is now certified and available on the new [Red Hat marketplace](#). The software platform has been tested and certified for the [Red Hat OpenShift container platform](#), the industry's most comprehensive enterprise Kubernetes platform, and can run anywhere OpenShift runs—on-premise or on any public or private cloud.

This offers many benefits:



ONE-CLICK INSTALLATION

With one click, you can install GigaSpaces InsightEdge on your OpenShift cluster, whether its on-premise or in the cloud. You indicate for each of your services how many servers and resources you want and their related scaling requirements, and everything is automatically provisioned for you. The operator manages not just the installation, but the entire lifecycle of the software, including automatic scaling, upgrades, and automatic recovery from failure. It also manages all hardware aspects, ensuring that data and its backup don't sit on the same physical machine.



DIGITAL SERVICES A/B TESTING

If you want to evaluate a new feature of a service, you can easily perform A/B testing. This is highly configurable; for example, you can direct 90% of your users to one service and 10% to another service as part of your test. This allows you to test new features and functionality within your applications on real users with no downtime in your production environment. Based on success metrics you define, you can decide whether to gradually move to the new version or roll back to the original.



RESOURCE MANAGEMENT AND ELASTICITY

Orchestrating GigaSpaces InsightEdge with Red Hat enables you to automatically scale up when you need more resources for peak loads and scaling out when you need more servers for intensive data processing. All you need to do is define your threshold and rules in the GigaSpaces platform. For example: If your CPU is at 70% utilization for more than two consecutive minutes, scale up by a factor of two. Resources are automatically scaled when that threshold is reached.



MONITORING AND MANAGING OF SERVICES

By leveraging OpenShift native controllers to monitor services, GigaSpaces InsightEdge can immediately detect and recover from software or hardware failures. In these cases, it signals the failure to the operational infrastructure and immediately re-provisions resources for the failed service. InsightEdge makes sure that data for each service is replicated on different servers to ensure it's always available, so the service can quickly recover without any downtime.



UPGRADES

Organizations can't afford to take mission-critical systems offline for maintenance. GigaSpaces InsightEdge leverages OpenShift capabilities to enable you to perform software upgrades yourself with no downtime.



SECURITY

GigaSpaces InsightEdge has been tested and certified for the Red Hat OpenShift container platform, widely considered to be significantly more secure than any other Kubernetes distribution.

Learn more about InsightEdge on the  Red Hat Marketplace

THE FUTURE IS BOTH CLOUDY & BRIGHT

Across all enterprises, [94% already use a cloud service](#). Multicloud adoption is growing as well; IDC says [70% of IT organizations are planning to implement strategic playbooks](#) for their multicloud approaches by 2023.

Digital transformation requirements are pushing more industries to migrate to hybrid, private, and public cloud infrastructures. As they do so, they are finding that they need to be able to leverage real-time insights from large datasets stored on disparate sources across multiple geographies. Thanks to a true cloud-native, in-memory data fabric, these enterprises are now proving able to transition from legacy systems to a modern architecture with minimal risk, enabling them to support growing volumes of data and an increased number of queries, while achieving speed, scale, availability, and high service levels, all in a unified environment.

Additional recommended reading:

Simplifying and Scaling your Hybrid and Multicloud Deployment Strategies with GigaSpaces v15.5
<https://www.gigaspaces.com/blog/simplifying-and-scaling-your-hybrid-and-multicloud-deployment-strategies-with-gigaspaces-v15-5/>

Rescaling Your Application in GigaSpaces
<https://docs.gigaspaces.com/latest/production/production-Rescaling-your-application-in-GigaSpaces.html>

Creating a Hybrid GigaSpaces Cluster – Cloud and Onprem
<https://docs.gigaspaces.com/latest/orchestration/gctl-create-gs-hybrid.html>

GigaSpaces InsightEdge is Now Available Through Red Hat Marketplace Worldwide
<https://www.gigaspaces.com/blog/gigaspaces-is-now-available-through-red-hat-marketplace/>



GigaSpaces is redefining in-memory technology to drive enterprise digital transformation with unparalleled speed, performance and scale. Never before has it been this simple to accelerate and scale real-time applications, analytics and operational BI on any data, at any load, across any environment.

The GigaSpaces InsightEdge Portfolio delivers the fastest and easiest to deploy suite of scalable software platforms to meet the most challenging enterprise data and analytics processing needs. The Smart Cache, Smart Operational Data Store (ODS) and Smart Augmented Transaction Platforms are powered by a unique combination of smart and autonomous data operations and management capabilities. One-click integrations with on-premise and cloud operational data stores, automatic data modeling, business-policy data tiering and dynamic scaling reduce time-to-market and ensure rapid response times and highest performance levels, with lower TCO.

Hundreds of Tier-1 and Fortune-listed organizations and OEMs across financial services, retail, transportation, telecom, and healthcare rely on GigaSpaces to optimize their business operations, enhance customer experience and comply with regulations.

GigaSpaces offices are located in the US, Europe and Israel with partners around the globe; serving customers such as Morgan Stanley, Bank of America, CSX, Goldman Sachs, Société Générale, Crédit Agricole, Avanza Bank, SITA, Avaya, Frequentis, CLSA, Groupe PSA and UBS.

This document is provided for information purposes only and the contents here of are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, or any purpose, without our prior written permission. GigaSpaces InsightEdge is a registered trademark of GigaSpaces and/or its affiliates. Other names may be trademarks of their respective owners.

