

IBM watsonx.data

Pricing Document



Pricing Document

IBM watsonx.data

IBM® watsonx.data® helps you access, integrate and understand all your data — structured and unstructured—across any environment. It optimises workloads for price and performance while enforcing consistent governance across sources, formats and teams.

SaaS (Consumption-Based)

- Priced per hour
- Select component size based on customer's workload characteristics.
- For pricing, since this is consumption-based, estimate the required size of individual components, determine the expected usage hours per month, and sum up all components, including the supporting service cost of 3 RU/hour.

Service	Part Number	List Price	Price Metric	Quantity*	Total Price
watsonx.data (License + S&S)	D0F4SZX	\$7,950 /core	VPCs (16 VPCs/node)	Ex) 160 VPCs	\$1,272,000
watsonx.data (S&S)	E0F4RZX	\$2,070 /core	VPCs (16 VPCs/node)	Ex) 160 VPCs	\$331,200
watsonx.data (Subscription)	D0F3HZX	\$3540 /core	VPCs (16 VPCs/node)	Ex) 160 VPCs	\$566,400
Expert Labs Services Bundle	D0GAVZX (Plan) D0GAUZX (Install) D0GATZX (Build)	\$19,00 \$17,200 \$17,200	Weeks	Qty 2	\$19,00 \$17,200 \$17,200
IBM watsonx.data as a Service Multicloud 1000 Resource Units (AWS only)	D0HLLZX (Rus) D0HLMZX (Overage)	\$1 Overage \$1.1	Resource Units	1000 RUs	Quantity dependent
IBM watsonx.data as a Service Consumption Multicloud Resource Unit Pay per Use (AWS only)	D0ZD7ZX	\$1 Overage \$1	Resource Units	1 RU	Quantity dependent
Watsonx As A Service Parts (IBM Cloud)	D0FHRZX WATSONX AS A SERVICE 1 US DOLLAR OVERAGE D0FHQZX WATSONX AS A SERVICE 1 US DOLLAR PER MONTH	\$1 Overage \$1	Resource Units	1 RU	Quantity dependent

watsonx.data Pricing & Packaging (Software)

Presto (java) Engine

Size	Entitlement	Compute	Memory	Description
Small	64 VPCs	64 vCPUs	512 GB	1 Coordinator node, 3 worker nodes. Each node assumes 16vCPU, 128 GB of memory, and 2 x 1900GB NVMe, starting from 10 GbE. Example of estimated workload supported by the configuration: 1 to 5 queries processing 100GB → 500 GB expanded/uncompressed active data in memory or up to 50 queries processing up to 8GB each. 100 - 200GB of Active data compressed on disk is required to be processed.
Medium	160 VPCs	160 vCPUs	1280 GB	1 Coordinator node, 9 worker nodes. Each node assumes 16vCPU, 128Gb of memory, and 2 x 1900GB NVMe, starting from 10 GbE. Example of estimated workload supported by the configuration: 1 to 5 queries processing 250GB → 1125 GB expanded/uncompressed active data. or up to 50 queries processing up to 22GB each. 250 - 500 GB Active data compressed on disk required to be processed.
Large	320 VPCs	320 vCPUs	2560 GB	1 Coordinator node 19 worker nodes. Each node assumes 16vCPU, 128Gb of memory, and 2 x 1900GB NVMe, starting from 10 GbE. Example of estimated workload supported by the configuration: 1 to 5 queries processing 500GB → 2250 GB expanded/uncompressed active data. or up to 50 queries processing up to 45GB each. 500 - 1000 GB of Active data compressed on disk required to be processed.
X-Large	1120 VPCs	1120 vCPUs	8960 GB	70 Nodes to configure as desired (i.e 10 clusters of 1 head node and 6 worker nodes or 1 large cluster with 70 nodes). Each node assumes 16vCPU, 128Gb of memory, and 2 x 1900GB NVMe, starting from 10 GbE. Example of estimated workload supported by the configuration: 1 to 5 queries processing 1750GB → 7875 GB expanded/uncompressed active data. or up to 175 queries processing up to 45GB each. 1750 - 3500 GB Active data compressed on disk required to be processed.
XX-Large	3200 VPCs	3200 vCPUs	25.6 TB	200 Nodes to configure as desired (i.e 10 clusters of 1 head node and 19 worker nodes or 1 large cluster with 200 nodes). Each node assumes 16vCPU, 128Gb of memory, and 2 x 1900GB NVMe, starting from 10 GbE. Example of estimated workload supported by the configuration: 1 to 5 queries processing 5000GB → 22,500 GB expanded/uncompressed active data or up to 500 queries processing up to 45GB each. 5000 - 10,000 GB Active data compressed on disk required to be processed.

watsonx.data Pricing & Packaging (Software)

Spark Engine

Size	Entitlement	Compute	Memory	Description
Small	6 VPCs	6 vCPUs	22 GB	1 Driver node, 2 executor nodes. The driver assumes 1 vCPU, and 4 Gb of memory. Each executor node assumes 1vCPU, and 4 Gb of memory. A Management Service occupies 3 vCPU and 10Gb of memory. Example of estimated workload supported by the configuration: 5 concurrent requests processing 10GB of data. Data Plane: 3 VPCs Small Control Plane: 3 VPCs Total VPCs: 6
Medium	24 VPCs	24 vCPUs	98 GB	1 Driver node, 5 executor nodes. The driver assumes 2 vCPU, and 12 Gb of memory. Each executor node assumes 3 vCPU, and 12 Gb of memory. A Management Service occupies 7 vCPU and 26Gb of memory. Example of estimated workload supported by the configuration: 15 concurrent requests processing 10GB of data. Data Plane: 17 VPCs Medium Control Plane: 7 VPCs Total VPCs: 24
Large	60 VPCs	60 vCPUs	242 GB	1 Driver node, 10 executor nodes. The driver assumes 3 vCPU, and 16 Gb of memory. Each executor node assumes 5 vCPU, and 20 Gb of memory. A Management Service occupies 7 vCPU and 26Gb of memory. Example of estimated workload supported by the configuration: 15 concurrent requests processing 10GB of data. Data Plane: 53 VPCs Medium Control Plane: 7 VPCs Total VPCs: 60

Note:

- Customers can have multiple Spark engines
- Sizing needs to be done based on the number of concurrent requests, data volume per request, and required driver/executor configurations, including management service overhead
- All watsonx.data components require watsonx.data Core



watsonx.data SaaS Pricing

- Consumption-based charges, driven off 3 Price metrics
- Quantity of nodes deployed
 - Duration of nodes deployed
 - Support services deployed (e.g., Access Control, [Metastore](#))

Consumption measured in Resource Units (RU). 1 RU = £1 GBP

Pricing	RU / Hour
Cache node cost per hour	£2.80
Compute node cost per hour	£6.50
Supporting services per hour	£3.00

Key Takeaways

- Usage can occur on IBM Cloud or AWS infrastructure
- Pricing depends on type of nodes, number of nodes used, and usage of nodes
 - AWS infrastructure provides two node types (Cache and Compute Optimized), IBM Cloud infrastructure has one node type (Cached Optimized). Prices are the same for AWS and IBM Cloud
 - Number of nodes relates to T-shirt sizes.
 - Usage relates to uptime of services.

Estimated T-Shirt Sizes for IBM Cloud

Sizing	Nodes	Total vCPU	Total RAM (GiB)	Resource Unit per Month*
Starter	2 nodes	32 vCPU	256 GB	•Using service 35% of the time: £3000
Small	4 nodes	64 vCPU	512 GB	•Using service 35% of the time: £4100
Medium	7 nodes	112 vCPU	896 GB	•Using service 35% of the time: £6000
Large	12 nodes	192 vCPU	1536 GB	•Using service 35% of the time: £9500

1 RU = 1 [watsonx](#) or [CPDaaS](#) credit = £1GBP

MCSP Quoting Details

Sub Id Dscr Full	Part Number	Part Dscr Long
IBM watsonx as a Service Multicloud Not Applicable	D0HLKZX	IBM watsonx as a Service Multicloud SVC Level Agreement
IBM watsonx.data as a Service Multicloud 1000 Resource Units	D0HLLZX	IBM watsonx.data as a Service Multicloud 1000 Resource Units per Annum
IBM watsonx.data as a Service Multicloud 1000 Resource Units	D0HLMZX	IBM watsonx.data as a Service Multicloud 1000 Resource Units Overage
IBM watsonx.data as a Service Multicloud 1000 Resource Units	E0HLNZX	IBM watsonx.data as a Service Multicloud 1000 Resource Advanced Support per Annum
IBM watsonx.data as Service Consumption Multicloud Resource Unit Pay per Use	D0ZD7ZX	IBM watsonx.data as Service Consumption Multicloud Resource Unit Pay per Use (AWS only)

IBM Cloud sometimes makes changes to offerings, included but not limited to: changing offering names, adding new offerings, removing existing offerings, adding or removing offering capabilities, changing costs per unit, changing units used to measure cost.

For further, updated information regarding all offerings, plus support and guidance on application running costs, contact your IBM representative or follow the link: <https://cloud.ibm.com/>

IBM Cloud provides a detailed cost estimator, and a down-to-the-penny cost calculator to help you plan for charges. You can check the actual cost after you build your apps through the IBM Cloud Dashboard. Where costs differ between this document and the IBM Cloud cost estimator tools, the tools should be considered to provide the correct price.

To calculate the cost of running an application on IBM Cloud, follow the link: <https://cloud.ibm.com/estimator>

There are the following billing options:

- Pay-As-You-Go: charged based on your use of provisioned Cloud compute and services
- Subscription: discounts are available based on a minimum monthly spending commitment and length of contract

IBM has a strategic MoU with Crown Commercial Services (CCS) which is applicable to all Public Sector Customers buying via G-Cloud. This provides agreed baseline levels of discounts and details higher levels of discounts for larger volumes or greater durations of contract. The IBM pages describing this are at https://ibm.biz/ccs_mou. Public Sector clients can request a full copy of the IBM MOU from CCS by sending an email to info@crownccommercial.gov.uk

The pricing in this G-Cloud Offering takes account of the MOU baseline discounts. As IBM's portfolio is somewhat complex and there are several factors that affect the pricing; if your requirements are for more than just one license for a month, we would suggest that you contact us directly to ask us to provide you with pricing with the correct levels of discount for the term and volumes that you require. The email address for G-Cloud related enquiries is ukcat@uk.ibm.com. Please ensure that you indicate which G-Cloud Offering(s) you are requesting details for.