

Urika®-GX Platform Technical Specifications

Seize your big data advantage with Cray's Urika®-GX agile analytics platform. The unmatched combination of system agility and pervasive speed delivers high frequency insights in days, not months.

PLATFORM CUSTOMIZATION OPTIONS

System Configurations	48, 32 or 16 nodes
Node Options	Processors: Intel® Xeon® E5-2600 v4 family 18-core or 8-core (2 per node) Memory: 128, 256 or 512 GB per node (32 GB or 16 GB DDR-2400 DIMMS, 8 or 16 per node) HDDs: 2 TB or 1 TB SATA 7.2K 2.5" 6 Gb/s (2 per node) SDDs: 0.8 TB, 1.6 TB or 4 TB (1 per compute node)

HARDWARE DETAILS: CONFIGURATION SPECIFICATIONS (STANDARD 48 NODES)

Rack	Single standard 42U/19" rack with 400mm rear extension for network cabling
Standard Options	18-core processor, 32 GB DIMM, 2 TB HDD, 0.8 TB SSD
Total Nodes	48 (44 compute, 2 I/O and 2 login)
Total Processors	96 (1,728 cores)
Total RAM	12 TB (256 GB per node)
Total Storage	227 TB total: 35 TB SSDs and 192 TB HDDs

COMPUTE NODE

Processor	36C (standard) 16C (optional)
RAM - DDR4-2133	256 GB (standard) 128 GB (optional) 512 GB (optional)
Hard Disk Drives (HDD)	4 TB (standard) 2 TB (optional)
Solid-State Drives	0.8 TB (standard) 1.6 TB (optional) 4 TB (optional)

LOGIN AND I/O NODES

Login Nodes	2 nodes supporting on-board GigE (optional 10 GigE or 40 GigE)
I/O Nodes	2 nodes supporting optional SAS, Fibre Channel, InfiniBand (FDR or EDR), 10 GigE or 40 GigE

Urika-GX Platform Technical Details

SOFTWARE DETAILS	
System Management	CentOS 7.2 Cray System Management Software based on OpenStack® technologies
Analytics Frameworks	<p>Hortonworks Data Platform</p> <ul style="list-style-type: none"> Includes HDFS, YARN, Hadoop®/MapReduce, Hive and HCatalog, ZooKeeper, WebHCat, Oozie, Pig, Mahout, Hue, Kafka, Flume and Sqoop <p>Apache Spark™</p> <ul style="list-style-type: none"> Spark Core, GraphX, MLLib, SparkR, Spark Streaming, Spark SQL, PySpark <p>CGE (Cray Graph Engine)</p> <ul style="list-style-type: none"> RDF triplestore, W3C standard SPARQL extended for mathematical algorithms <p>Resource Management Ecosystem</p> <ul style="list-style-type: none"> Apache Mesos™ Marathon <p>Analytics Programming Environment</p> <ul style="list-style-type: none"> OpenJDK, Scala, R, CPython, Anaconda Python, Intel MKL, NumPy, SciPy, Maven, Scala Build Tool (sbt) <p>Additional Analytics Software</p> <ul style="list-style-type: none"> Jupyter Notebook (JupyterHub) configured for Python, R, Spark, PySpark, SparkR
External File Systems	Cray Sonexion® storage (Lustre®), GPFS and NFS connectivity, or other POSIX-compliant global storage (optional)