

## 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	
<b>System Configurations</b>	48, 32 or 16 nodes
<b>Node Options</b>	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)	
<b>Rack</b>	Single standard 42U/19" rack with 400mm rear extension for network cabling
<b>Standard Options</b>	18-core processor, 32 GB DIMM, 2 TB HDD, 0.8 TB SSD
<b>Total Nodes</b>	48 (44 compute, 2 I/O and 2 login)
<b>Total Processors</b>	96 (1,728 cores)
<b>Total RAM</b>	12 TB (256 GB per node)
<b>Total Storage</b>	227 TB total: 35 TB SSDs and 192 TB HDDs
COMPUTE NODE	
<b>Processor</b>	36C (standard) 16C (optional)
<b>RAM - DDR4-2133</b>	256 GB (standard) 128 GB (optional) 512 GB (optional)
<b>Hard Disk Drives (HDD)</b>	4 TB (standard) 2 TB (optional)
<b>Solid-State Drives</b>	0.8 TB (standard) 1.6 TB (optional) 4 TB (optional)
LOGIN AND I/O NODES	
<b>Login Nodes</b>	2 nodes supporting on-board GigE (optional 10 GigE or 40 GigE)
<b>I/O Nodes</b>	2 nodes supporting optional SAS, Fibre Channel, InfiniBand (FDR or EDR), 10 GigE or 40 GigE

## Urika-GX Platform Technical Details

### SOFTWARE DETAILS

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