



Urika-GX System Overview

(2.2.UP00)

S-3017

Contents

1 About the Urika®-GX System Overview Publication.....	3
2 Features of Urika-GX High Performance Analytic Platform	5
3 Urika-GX Hardware Components.....	9
3.1 Network Components.....	10
4 Urika-GX Software Components.....	13
4.1 File Systems.....	14
4.2 Application and Resource Management Tools.....	14
4.3 System Management and Monitoring Tools.....	15
5 Urika-GX Service Modes.....	18

1 About the Urika®-GX System Overview Publication

This publication describes the features, hardware and software of the Cray® Urika®-GX system. It also provides an overview of the analytic applications and third party tools installed on Urika-GX.

Typographic Conventions

<i>Monospace</i>	Indicates program code, reserved words, library functions, command-line prompts, screen output, file/path names, key strokes (e.g., <code>Enter</code> and <code>Alt-Ctrl-F</code>), and other software constructs.
Monospaced Bold	Indicates commands that must be entered on a command line or in response to an interactive prompt.
<i>Oblique or Italics</i>	Indicates user-supplied values in commands or syntax definitions.
Proportional Bold	Indicates a graphical user interface window or element.
\ (backslash)	At the end of a command line, indicates the Linux® shell line continuation character (lines joined by a backslash are parsed as a single line). Do not type anything after the backslash or the continuation feature will not work correctly.

Scope and Audience

The audience of this publication is users and system administrators of the Urika®-GX system. This publication is not intended to provide detailed information about open source products used in the system. References to online documentation are included where applicable.

Record of Revision

Date	Addressed Release
September, 2018	2.2UP00
March, 2018	2.1UP00
December, 2017	2.0UP00
April, 2017	1.2UP00
December, 2016	1.1UP00
August, 2016	1.0UP00
March, 2016	0.5UP00

Record of Revision

This revision contains editorial changes.

Trademarks

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, Urika-GX, Urika-XA, Urika-GD, and YARCDATA. The following are trademarks of Cray Inc.: APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYDOC, CRAYPAT, CRAYPORT, DATAWARP, ECOPHLEX, LIBSCI, NODEKARE. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.

2 Features of Urika-GX High Performance Analytic Platform

The Cray® Urika®-GX system is a big data platform optimized for analytic work-flows. It combines a highly advanced hardware platform with a comprehensive analytic software stack to help derive optimal business value from data. The Urika-GX platform provides the tools required for capturing, analyzing and organizing a wide variety of data types from different sources.

Major features of the Urika-GX system include:

- **Multi Tenancy** - On Urika-GX, multi tenancy is implemented through the use of tenant VMs that run on physical nodes and provide controlled access to services on the physical nodes through a command proxy mechanism. A number of CLI scripts can be used to perform tenant management tasks.

Urika-GX enables:

- using Spark, HDFS, Lustre and NFS in a secure multi tenant environment
- running containerized services, which are managed via Kubernetes
- isolation of Spark jobs, data, network and user activity between individual tenant VMs
- customization of application configuration for individual tenant VMs

For more information, refer to "*Urika-GX System Administration Guide*".

- **System, application and resource management tools**
 - Application and resource management tools - There are a number of resource and application management tools installed on the Urika-GX system, including Mesos, Marathon YARN, `mr-un` and Kubernetes. These tools enable management of analytic workloads, dynamically allocate system resources to applications as needed, and provide the flexibility of running multiple jobs across the cluster concurrently. For increased system reliability, Urika-GX features containerization of Spark applications. Containers are managed via Kubernetes for optimized system reliability. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".
 - Performance analysis tools - Urika-GX features a set of performance analysis tools to analyze performance of Urika-GX's analytic applications. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".
 - Urika-GX CLI - Urika-GX features a number of CLI commands that enable monitoring and managing analytic services. These commands are also integrated in to the system management UI for ease of use.
 - System Monitoring tools - The system features Nagios® Core for monitoring system resources, such as usage of CPU, memory, network bandwidth, Lustre, HDD and SSD storage. It also enables monitoring node status. For more information, see S-3016, "*Urika®-GX System Administration Guide*".
- **Cray Programming Environment (PE)** - Urika-GX features support for using the Cray PE, which comprises of software components required for developing, compiling, executing, debugging, and analyzing code.
- **Analytic application development tools**
 - Urika-GX analytic programming environment - Urika-GX features an analytic programming environment that comprises of a set of tools that facilitate building, debugging, and compiling various types of analytic

applications using various programming languages, including Python, R, Java, Spark etc. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".

- Jupyter Notebook - Urika-GX ships pre-installed with the Jupyter notebook, which enables creating executable documents. In addition, there are some cookbooks that Urika-GX ships with that provide information about using the system. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".
- Cray Graph Engine (CGE) - Urika-GX also features CGE for performing complex graph analytic tasks on RDF databases at interactive speeds. For more information, refer to S-3014, "*Cray™ Graph Engine User Guide*" and S-2591, "*Cray™ Graph Engine Quick Start Guide*."
- Cray Application Management UI- Urika-GX features a user friendly UI for monitoring and managing jobs submitted to the system. In addition, the system features the Urika®-GX Applications Interface, which provides access to various system and analytic applications installed on the system. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*" and S-3016, "*Urika®-GX System Administration Guide*".
- Additional analytic tools
 - Urika-GX system features a number of data analytic tools that help perform analytic tasks, such as executing Hadoop and SPARK jobs. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".
 - `mrunc` is a Cray-built application launcher program that enables running parallel jobs on Urika-GX using resources managed by Mesos/Marathon. In addition to launching jobs, `mrunc` also has the ability to display currently active Mesos Frameworks and Marathon applications, provide extensive details on running Marathon applications, as well as cancelling/ stopping currently active Marathon applications. CGE uses `mrunc` to launch jobs under the Marathon framework on the Urika-GX system.
- **Security** - Urika-GX features tight application and data security for certain applications. The system provides two security modes, 'Secure' and 'Default', which dictate the list of applications and services that can be used. For more information, refer to "*Urika-GX System Administration Guide*".
- **Data virtualization tools** - Urika-GX features the Cray Data Virtualization Service (DVS), which is a distributed network service that projects local file systems resident on I/O nodes or remote file servers to compute and service nodes within the Cray system. DVS-specific configuration settings enable clients (compute nodes) to access a file system projected by DVS servers. Thus, Cray DVS represents a software layer that provides scalable transport for file system services. For more information, see S-3016, "*Urika®-GX System Administration Guide*".
- **Data analysis and visualization tools** - Urika-GX features connectivity to Tableau® for efficient data visualization of analytics work-flow processing. Tableau connects easily to Urika-GX, as well as to other data sources. It enables transforming Urika-GX data into interactive, visually appealing visualizations called dashboards, quickly and easily. For more information, visit <https://www.tableau.com> and refer to S-3015, "*Urika®-GX Administration Guide*".
- **Support for external storage** - Urika-GX systems feature support for external POSIX-compliant storage systems. Currently, Lustre® is supported as an external file system.

If the Cray Sonexion system is being used as an external storage system, refer to "*Sonexion Administrator Guide*" at <https://pubs.cray.com>
- **Containerization** - Urika-GX features the infrastructure for running Docker containers orchestrated by Marathon and Mesos, pulled from the public Docker Hub repository. Docker enables packaging applications with their dependencies without having to worry about system configurations, thus making the applications more portable.
- **Multihoming** - Urika-GX enables users to transfer data directly from external sources to the compute nodes of Urika-GX, as well as read data stored on Urika-GX via supported APIs from an external application. So, users can access data streaming applications and services directly from compute nodes via the Urika-GX

operational network, resulting in optimized performance (1Gb/sec/node) and network utilization. Users can also access application log files residing on the compute nodes where the tasks belonging to the application run.

- **Fault tolerance** - A number of Urika-GX's analytic applications are optimized for system resiliency and fault tolerance. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".
- **Hardware** - Urika-GX rack and sub-rack enclosures use GreenBlade™ technology, which provides reliable and high-quality blade-based solutions. GreenBlade's unique design decreases power consumption per node, thereby significantly reducing data center energy costs. For more information, see "*Urika®-GX Hardware Guide*".
- **Flexible configurations** - The Urika-GX rack can support up to 3 sub-rack chassis, each of which contains 16 nodes. The Urika-GX system contains three types of nodes:
 - Compute nodes - Compute nodes run application programs and services.
 - I/O nodes - I/O nodes facilitate connecting to external storage and file systems. Urika-GX currently supports Lustre as an external file system.
 - Login nodes - Login nodes enable users to log on to the Urika-GX system and facilitate launching jobs from the command line.

The Urika-GX system contains 2 login and 2 I/O nodes, whereas the remaining nodes in the system are all compute nodes.

- Node options
 - Processor: Intel® Broadwell 18C E5-2697 v4 or 8C E5-2620 v4 (2 per node)
 - Memory: 128, 256 or 512 GB per node (32 GB or 16 GB DDR-2400 DIMMS/DDR-2133 DIMMs) (8 or 16 per node)
 - HDDs: 2 TB or 1 TB SATA 7.2K 2.5" 6 Gb/s (2 per node) processors
 - SSDs: 2 TB or 4 TB Intel® PCIe (1 per compute node)
- Connectivity options
 - I/O nodes - Based on the external storage, the following PCIe HBA options are available for I/O nodes (only one can be selected out of the following):
 - LSI Logic LSI00343 9300-8E SGL 8-Ports Ext 12GB/S SATA SAS PCIe3.0
 - Mellanox MCX354A-QCBT ConnectX-3 VPI Adapter Card Dual-Port QSFP QDR IB (40Gb/s) and 10GbE PCIe3.0 x8
 - Mellanox MCX354A-FCBT ConnectX-3 VPI Adapter Card Dual-Port QSFP FDR IB (56Gb/s) and 40/56GbE PCIe3.0 x8
 - ConnectX-4 VPI cards
 - Mellanox MCX455A-ECAT ConnectX-4 VPI adapter card, EDR IB (100Gb/s) and 100GbE, single-port QSFP, PCIe3.0 x16
 - Mellanox MCX456A-ECAT ConnectX-4 VPI adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP, PCIe3.0 x16
 - QLogic® QLE2672 Gen 5 (16Gb) Fibre Channel (FC) Adapter

NOTE: Only optical fibre channel cables should be used with Urika-GX. The QLE2672 includes optical SFPs. However, when using copper fibre channel cables, the optical SFPs are removed because copper fibre channel cables include their own integrated copper SFP transceiver. But the mechanical locking of the copper cable's SFP requires a tool to release once inserted into the QLE2672 HBA. Due to the close proximity of the

Urika-GX node blades, there is lack of sufficient room to utilize a tool to release the copper cable SFP. Therefore, only optical fibre channel cables are supported.

- Login nodes - Login nodes support on-board GigE connectivity by default. A PCIe card supporting 10GigE connectivity can optionally be added to a login node for connectivity to an external network.
- Processor options (compute nodes) - Intel® Broadwell 18C E5-2697 v4 or 8C E5-2620 v4 (2 per node)
- Hard disk drive options (compute nodes)
 - 2TB SATA storage per node provided by 2 x 1TB SATA 7.2K 2.5" 6GB/S HDDs
 - 4TB SATA storage per node provided by 2 x 2TB SATA 7.2K 2.5" 6GB/S HDDs. This is the default configuration.
- Memory options (compute nodes)
 - 128GB per node provided by 8 x 16GB DDR4-2133 DIMMs
 - 256GB per node provided by 8 x 32GB DDR4-2133 DIMMs. This is the default configuration.
 - 512GB per node provided by 16 x 32GB DDR4-2133 DIMMs
 - DDR4 2400MHz DIMM availability with broadwell processors

All compute nodes include PCIe 3.0 SSD storage and are configured identically in terms of processor, memory, HDD, and SSD storage. I/O nodes do not support SSDs. Login nodes do not include SSD storage by default and may optionally include a PCIe SSD. If an SSD is included, it needs to be the same SSD that is installed on compute nodes. If an SSD is not needed on the login node, network bandwidth can be improved by installing one of the following optional PCIe network adaptors:

- Mellanox MCX312A-XCBT ConnectX-3 Network Interface Card Dual-Port 10GbE PCIe3.0 x8
- Mellanox MCX354A-QCBT ConnectX-3 VPI Adapter Card Dual-Port QSFP QDR IB (40Gb/s) and 10GbE PCIe3.0 x8
- Mellanox MCX354A-FCBT ConnectX-3 VPI Adapter Card Dual-Port QSFP FDR IB (56Gb/s) and 40/56GbE PCIe3.0 x8
- Mellanox MCX455A-ECAT ConnectX-4 VPI adapter card, EDR IB (100Gb/s) and 100GbE, single-port QSFP, PCIe3.0 x16
- Mellanox MCX456A-ECAT ConnectX-4 VPI adapter card, EDR IB (100Gb/s) and 100GbE, dual-port QSFP, PCIe3.0 x16
- Mellanox MCX415A-CCAT ConnectX-4 Network Interface Card, 100GBE single-port QSFP, PCIe3.0 X16
- Mellanox MCX416A-CCAT ConnectX-4 Network Interface Card, 100GBE dual-port QSFP, PCIe3.0 X16

3 Urika-GX Hardware Components

Major hardware components of the Urika®-GX system rack include:

Sub-rack Enclosure

There are up to three GreenBlade 10.5 Rack Unit (RU) SR10216 sub-rack enclosures contained within the Urika-GX system rack.

Each Urika-GX sub-rack contains the following components:

- **Nodes** - The Urika-GX system rack supports up to three sub-racks, each of which contains 16 GreenBlade™ GB522XA blades/nodes. All Urika-GX nodes run the CentOS 7.2 operating system and portions of the Cray Linux Environment (CLE).
- **Aries Network Card Controller (ANCC)** - Each sub-rack chassis of the Urika-GX rack contains two Dual Aries Network Card (dANC) boards. Each dANC board in turn contains 2 Aries Network Card Controllers (ANCCs), an Advanced RISC Machines (ARM) processor, and a number of environmental sensors. The ANCCs support High Speed Network (HSN) connectivity to all the nodes.
- **Intelligent Server Control Board (iSCB)** - Each sub-rack contains one iSCB module, which reports the system's health status. The iSCB module interfaces with the Hardware Supervisory System (HSS), which is used for system monitoring and management.

In addition to the aforementioned components, each Urika-GX sub-rack contains 6 Power Supply Units (PSUs) and 6 cooling/fan modules.

System Management Workstation (SMW)

The SMW is a Dell R630 server that acts as a single-point interface to a system administrator's environment. The SMW provides a terminal window for performing administrative tasks, such as adding user accounts, changing passwords, and monitoring nodes and applications.

The SMW contains an Intel Xeon E5-2623 v4 2.6 GHz processor, 64GB of RAM and 5TB of physical storage. The base operating system of the SMW is CentOS 7.3

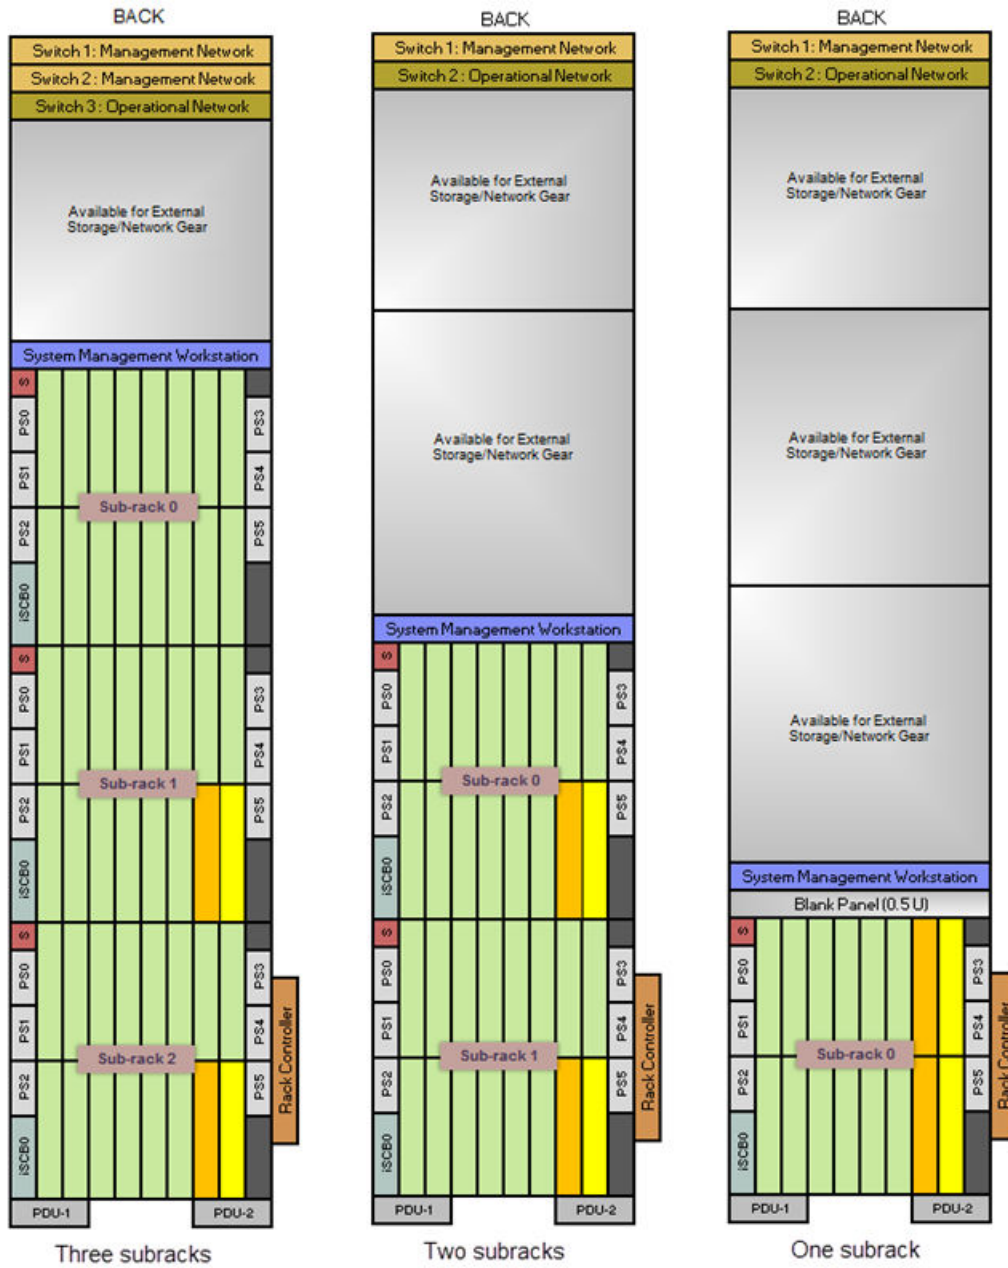
Switches

The Urika-GX rack contains three single RU 48-port GigE Brocade ICX 6450 Ethernet switches for management and external connectivity.

Rack Controller (RC)

The RC acts as a component of the HSS and manages communication between the SMW and the rack, sub-rack, dANCCs, and, in some cases, nodes. It monitors the sub-rack's temperature and other external environmental factors and is installed on the front interior of the rack side panel.

Figure 1. Urika-GX System Rack Configurations



Nodes within each subrack

- = Compute node
- = Log-in node
- = I/O node

PDU = Power distribution unit

PS = Power supply

iSCB = Intelligent sub-rack control board

S = iSCB DIP switch

3.1 Network Components

There are 3 networks deployed on the Urika®-GX platform:

- **Aries High Speed Network (HSN)** - The Aries HSN provides high speed application and data network connectivity between nodes. This network provides node interconnect via high bandwidth, low latency DMA access. The hardware to support this network consists of an Aries Interface Board (AIB) connected to an available PCIe slot on each Urika-GX node and integrated into the node chassis assembly. The AIB is connected to the dANC integrated in the Urika-GX sub-rack. Copper cables provide an all-to-all connection of all dANCs in the system.
- **Operational Ethernet network**- The operational Ethernet network is used for ingesting user data. This network is comprised of a single unit 48-port GigE switch that provides dual 1GigE and/or dual 10GigE interfaces to the site network. Urika-GX's login nodes do not route through this switch and need to be directly connected to the site network. The operational network allows node connectivity externally from Urika-GX to the site network. The Urika-GX compute and I/O nodes are connected to a single managed Brocade ICX 6450-48, 48 port switch with a single power supply. Connectivity of this network to the site network is made possible by two available Gigabit Ethernet ports and/or two 10 Gigabit Ethernet ports on the ICX 6450-48 switch.

The operational network can also be used to access data streaming applications and services directly from compute nodes.

- **Management Ethernet network** - The management Ethernet network is primarily used for system management, and not for user data. The management Ethernet network is comprised of two stacked 1U 48-port switches, which are located at the top of the Urika-GX rack, and can optionally contain redundant switch power supplies. These switches provide GigE management Ethernet connectivity to every node, System Management Workstation (SMW), Rack Controller (RC), Intelligent Subrack Control Board (iSCB), Power Distribution Units (PDUs), Dual Aries Network Cards (dANCs) and to the operational network that connects to the nodes.

The Urika-GX system also contains the following subnets:

- SMW subnet, which provides connectivity to the SMW and the RC.
- Rack subnet, which provides connectivity to the dANCs and iSCB module.

This network is supported by two managed Brocade ICX 6450-48, 48 port switches stacked together with two 10GigE optical interconnects. Each switch contains a single power supply, and can optionally contain redundant switch power supplies. The following VLANs are defined for this network to support management network traffic:

- VLAN 102 - Uses ports 1-5 on each ICX 6450-48 switch. This is a dual-mode (tagged dual-mode for VLAN 102 and tagged for VLAN 103) VLAN. Untagged traffic on these ports belongs to VLAN 102. Traffic can be tagged for VLAN 103. The SMW HSS interface, the RC for a given rack, and the PDUs for a given rack are connected to these ports.
- VLAN 103 Ports 6-12 on each ICX 6450-48 switch. Untagged traffic on these ports belongs to VLAN 103. The iSCBs and dANC cards are connected to these ports.
- VLAN 104 Ports 13-48 on each ICX 6450-48 switch.

NOTE: Traffic on this VLAN may be reduced if VLAN 105 is needed for storage as long as each compute node is connected to VLAN 104

Untagged traffic on these ports belongs to VLAN 104. The compute nodes and the SMW node-side network are connected to these ports.

- VLAN 105 Some number of Ports 13-48 on each ICX 6450-48 switch, as needed for storage management. Untagged traffic on these ports belongs to VLAN 105. The Storage Management Ports are connected to these ports.

- VLAN 1 (default) is unused.

Traffic from the SMW to the subcomponents in the rack subnet, and vice versa, is routed through the corresponding RC.

For additional information, see the *Urika®-GX Hardware Guide*.

4 Urika-GX Software Components

The Urika®-GX system runs a combination of Cray-developed, third-party, and open-source software components. The system software is optimized for applications that have fine-grain synchronization requirements, large processor counts, and significant communication requirements.

Major software components of Urika-GX are listed below:

Operating System

Urika-GX compute nodes and SMW run the CentOS 7.4 operating system. Along with CentOS, Urika-GX nodes also run portions of the Cray Linux Environment (CLE) for providing a number of Cray-specific features and functionalities.

Analytics Applications

- Cray Graph Engine (CGE) - CGE is a highly optimized software application designed for high-speed processing of interconnected data. It is a SPARQL database engine that features an advanced platform for searching very large, graph-oriented databases and querying for complex relationships between data items with a database. It provides the tools required for capturing, organizing and analyzing large sets of interconnected data. CGE enables performing real-time analytics on the largest and most complex graph problems, and features highly optimized support for inference, deep graph analysis, and pattern-based queries. For more information, see S-3014, "*Cray® Graph Engine (CGE) User Guide*".
- Spark and Hadoop core and ecosystem components - Urika-GX uses the Hortonworks Data Platform (HDP) for Hadoop support. In addition to the core Hadoop and Spark components, Urika-GX ships with a number of ecosystem components for increased productivity.
- Urika-GX analytics programming environment - Urika-GX features an analytic programming environment that comprises of a set of tools that facilitate building, debugging, and compiling various types of analytics applications using various programming languages, including Python, R, JAVA, Spark etc. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".

System Management Tools

The Urika-GX system features the Cray HSS for system management. HSS software runs on the SMW.

Cray Programming Environment

To facilitate running High Performance Computing (HPC) jobs using the Cray programming environment, Urika-GX features The Cray compiler, Cray Message Passing libraries, Cray scientific libraries and Cray performance tools along with Cray modules to easily load them into the environment. These tools help in developing, compiling, executing and analyzing code.

Additional Third Party Software

Urika-GX also features a number additional tools, such as the Jupyter Notebook and Docker. Urika-GX supports connectivity with Tableau® for efficient data visualization of analytics work-flow processing.

4.1 File Systems

Supported file system types on Urika-GX include:

- **Internal file systems**

- Hadoop Distributed File System (HDFS) - Hadoop uses HDFS for storing data. HDFS is highly fault-tolerant, provides high throughput access to application data, and is suitable for applications that have large data sets. Urika-GX also features tiered HDFS storage. HDFS data is transferred over the Aries network.
- Network File System (NFS) - The Urika-GX SMW hosts NFS, which is made available to every node via the management network.
- `/mnt/lustre` - This is a directory that hosts Lustre file system data if DAL/Sonexion is used.



CAUTION: Avoid using NFS for high data transfers and/or large writes as this will cause the network to operate much slower or timeout. NFS, as configured for Urika-GX home directories, is not capable of handling large parallel writes from multiple nodes without data loss. Though It is possible to configure NFS to handle parallel writes, it would require a hard mount, which would have undesired consequences.

File Locations

- Home directories are mounted on (internal) NFS, with limited space
 - Distributed file system (Lustre), if provisioned, is mounted at `/mnt/lustre` and is suitable for larger files.
- Lustre mounts are isolated, with individual tenants having their own mount point.

4.2 Application and Resource Management Tools

There are a number of application and resource management tools used on the system, depending on the security mode the system is running under. These include:

Application and Resource Management Tools Used Under Each Security Mode

- **Default Mode**

- Apache™ Mesos™ - Apache Mesos simplifies the process of deploying and managing applications in large-scaled clustered settings. It can run multiple applications on a dynamically shared pool of nodes. Mesos consists of a master daemon that manages agent daemons running on each cluster node, and Mesos applications (also called frameworks) that run tasks on the agent daemons. The master enables fine-grained sharing of resources across applications by making them resource offers and decides the number of resources to offer each framework according to a given organizational policy.

- Apache™ Yet Another Resource Negotiator (YARN) - YARN is a cluster management tool, responsible for managing compute resources in a cluster and using those resources for scheduling of applications. YARN is designed to allow multiple, disparate applications to run on a multi-tenant platform. It is responsible for managing and monitoring workloads, maintaining a multi-tenant environment, implementing security controls, and managing high availability features of Hadoop. YARN acts as the resource manager for Hadoop jobs on the Urika-GX system.
- `mr`un - Mesos acts as the primary resource manager on Urika-GX when the system is running in the default security mode. As part of the Mesos ecosystem, Urika-GX features the Cray-developed `mr`un application launcher program. `mr`un enables running parallel jobs on Urika-GX, using resources managed by Mesos/Marathon. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".

When the system is running in the default security mode:

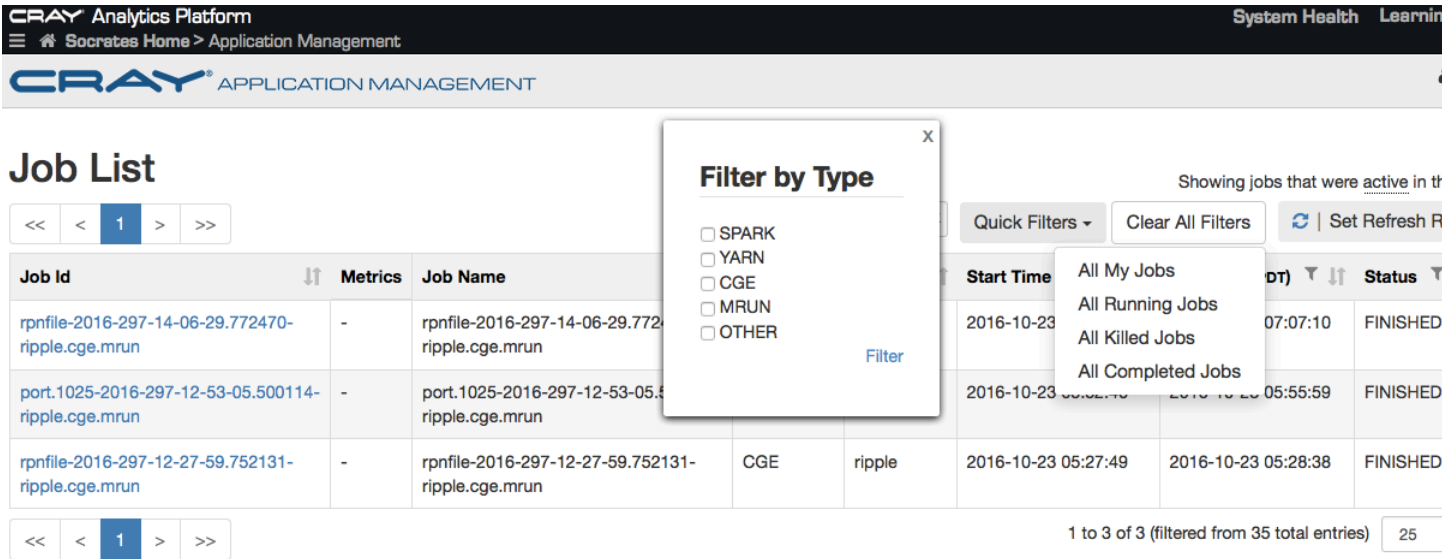
- Spark jobs run directly on Mesos
- Hadoop jobs run on YARN, which runs on Marathon, which in turn runs on Mesos
- CGE jobs run on `mr`un, which runs on Marathon, which in turn runs on Mesos
- **Secure Mode**
 - **Kubernetes** - Kubernetes is used for automating the deployment, scaling, and management of containerized applications. It groups containers that make up an application into logical units for simplified management and discovery. On Urika-GX, Kubernetes is used under the secure mode to manage Spark and HDFS jobs.

4.3 System Management and Monitoring Tools

There are a number of tools that can be used for managing and monitoring the Urika®-GX system. These tools include:

- **Hardware Supervisory System (HSS)** - HSS is an integrated system of hardware and software components that monitors components, manages hardware and software failures, controls system start up and shutdown, manages the system interconnection network and maintains system states. The HSS communicates with nodes and management processors over an internal (private) Ethernet network that operates independently of the system interconnection network.
- **Nagios® Core** - Nagios Core is a system monitoring tool that enables monitoring node status as well as network bandwidth, Lustre, SSD, HDD, CPU and memory usage.
- **Cray Application Management** - Urika-GX features the Cray Application Management UI that enables viewing, searching, filtering and performing actions on jobs submitted to the Urika-GX system

Figure 2. Cray Application Management UI



- **Analytic applications monitoring tools** - There are a number of tools installed on the Urika-GX system that facilitate monitoring the analytic applications. These tools include Grafana, Hadoop Job History Server, and Spark History Server. For more information, see S-3015, "*Urika®-GX Analytic Applications Guide*".

Urika-GX features the **Urika®-GX Applications Interface**, which enables accessing many of the pre-installed analytic applications, the system management UI and Urika-GX documentation.

Figure 3. Urika®-GX Applications Interface UI



5 Urika-GX Service Modes

Urika-GX features application and data security for a number of applications. Many security mechanisms play a role in the overall security architecture to ensure the system is protected against unauthorized access.

- **Default Mode** - All the installed analytic applications are available under the default service mode. Applications are configured with basic security levels wherever possible, though certain applications may provide no security features. More specifically, HDFS runs using its simple authentication mode under the default service mode.
- **Secure Mode** - Urika-GX uses Kerberos authentication while running under the secure mode. Only a limited number of analytic applications are available under this mode and those applications are configured to run in a secure configuration, the exact details of which vary by application. Typically, this means that applications that interact with HDFS require valid Kerberos credentials. Moreover, user interfaces are either disabled or require authentication under this mode.



CAUTION: If the Urika GX system is running in the secure mode, Cray does not recommend toggling back to the default mode while in production. In the default service mode, the security assurances provided by secure service mode are not in place and the security of data that was protected by secure mode may be compromised while running in the default mode. Cray cannot extend the secure mode security assurances to any system that has run in a production state in the default mode until that system has been fully re-deployed.

Table 1. List of Services Available Under the Default and Secure Service Modes

Service	Available in Default Service Mode	Available in Secure Service Mode
Cray Programming Environment	Yes	No
SELinux	Yes	No
Analytic Applications and Resource Management Tools		
ZooKeeper	Yes	Yes
Spark	Yes	Yes
Mesos Master	Yes	No
Mesos Slave	Yes	No
Kubernetes	No	Yes
HDFS NameNode	Yes	Yes
HDFS Secondary NameNode	Yes	Yes
HDFS DataNode	Yes	Yes
CGE	Yes	No

Spark History Server	Yes	No
Spark Thrift Server	Yes	No
YARN	Yes	No
YARN Resource Manager	Yes	No
YARN Node Managers	Yes	No
Hadoop Job History Server	Yes	No
Hadoop Application Timeline Server	Yes	No
Hive MetaStore	Yes	No
HiveServer2	Yes	No
Hive WebHCat	Yes	No
HUE	Yes	No
Oozie Server	Yes	No
Marathon	Yes	No
Grafana	Yes	No
InfluxDB	Yes	No
JupyterHub	Yes	No
System Management Tools		
HSS	Yes	Yes
Nagios	Yes	Yes
Cobbler	Yes	Yes
Secret Manager	Yes	Yes
Tenant management and proxy tools	Yes	Yes
Analytic programming environment components <ul style="list-style-type: none"> • R • Numpy • Scipy • GIT • Environment modules • glibc-devel • gcc • Python 34 • Python 27 • Scala • Apache Maven 	Yes	Yes

• Anaconda Python		
Miscellaneous Tools and UIs		
YAM	Yes	No
mrunc	Yes	No
Urika Application Management (UAM)	Yes	No
Urika-GX Application Interface (UAI)	Yes	No
HAProxy	Yes	Yes
Connectivity to Tableau	Yes	No
Docker	No	Yes

Any additional services installed on the system will use their own security mechanisms and will not be affected by Urika-GX's default and secure modes.

Table 2. Relationship Between Access Levels and Service Modes

Mode	Restricted Access		Relaxed Access	
	As Member on Tenant	On Physical Node	As Member on Tenant	On Physical Node
Secure	Has proxied access to Spark and HDFS commands	No access	Has proxied access to Spark and HDFS commands	Has direct access to all the services supported in the secure mode
Default	Unsupported	Unsupported	Unsupported	Has direct access to all the services supported in the default mode