



# **Cray® Graph Engine User Guide (2.0.UP00)**

## **S-3014**

# Contents

1 About the Cray® Graph Engine User Guide.....	5
2 About the Cray Graph Engine (CGE).....	7
2.1 Cray Graph Engine (CGE) Features.....	7
2.2 Concepts of Operation.....	7
2.2.1 What the Cray Graph Engine (CGE) is Not: a Relational Database.....	7
2.2.2 What the Cray Graph Engine (CGE) is: a Graph-Oriented Database that Uses RDF.....	8
2.3 About SPARQL .....	10
2.4 System Architecture Overview.....	10
2.5 RDF and SPARQL Resources.....	11
3 Building a Database .....	12
3.1 About Rules Files .....	14
3.2 Create a Set of Inferencing Rules .....	14
3.3 Sample RDFS Rules File .....	17
3.4 Limitations to Jena Rules Syntax .....	19
4 Launch the CGE Server Using the <code>cge-launch</code> Command.....	20
5 Mechanisms to Interact with the Cray Graph Engine (CGE) Database.....	22
5.1 Cray Graph Engine (CGE) Command Line Interface.....	22
5.1.1 Cray Graph Engine (CGE) Command Output.....	23
5.1.2 Cray Graph Engine (CGE) CLI Common Options.....	24
5.1.3 SSH Identities.....	28
5.1.4 CGE Hadoop HDFS Configuration.....	28
5.1.5 Cray Graph Engine (CGE) Properties File.....	29
5.1.6 Create Checkpoints Using the CGE <code>checkpoint</code> Command.....	31
5.1.7 Compile SPARQL Commands Using the CGE <code>compile</code> Command.....	32
5.1.8 Check the Database State Using the CGE <code>echo</code> Command.....	33
5.1.9 Launch the CGE Web Server Using the <code>fe</code> Command.....	34
5.1.10 Determine How Locations Are Being Searched Using the <code>get-configuration</code> Command.....	34
5.1.11 Display keyword ID and User Friendly Keyword Name Mappings Using the <code>keyword-</code> <code>lookup</code> Command.....	35
5.1.12 Retrieve Default Server Logging Information Using the <code>log-info</code> Command.....	35
5.1.13 Lookup Mappings Between Log level Values and User Friendly Log Level Names Using the <code>log-lookup</code> Command.....	36
5.1.14 Change the Default Logging Configuration of the CGE Server Using the <code>log-</code> <code>reconfigure</code> Command.....	36

5.1.15 Display Information About the Default NVP Configurations Using the CGE <code>nvp-info</code> Command.....	37
5.1.16 Change Default NVP Configurations Using the CGE <code>nvp-reconfigure</code> Command.....	37
5.1.17 Display Server Output Directory Information Using the <code>output-info</code> Command .....	37
5.1.18 Change the Server's Output Directory Using the CGE <code>output-reconfigure</code> Command .....	38
5.1.19 Execute Queries Using the CGE <code>query</code> Command.....	38
5.1.20 Shutdown the CGE Server Using the <code>shutdown</code> Command.....	41
5.1.21 Execute Sparql Queries and Updates Using the <code>sparql</code> Command.....	41
5.1.22 Execute Updates on a Database Using the CGE <code>update</code> Command.....	42
5.2 Access the Cray Graph Engine (CGE) Graphical User Interface .....	43
5.2.1 Execute SPARQL Queries Using the CGE UI.....	45
5.2.2 Execute SPARQL Updates Using the CGE Update Interface.....	45
5.2.3 Create a Checkpoint Using the CGE UI.....	46
5.2.4 Cray Graph Engine (CGE) Advanced Options .....	47
5.2.5 View Server Configurations Using the CGE UI.....	48
5.2.6 Edit Server Configurations Using the CGE UI.....	49
5.3 SPARQL Endpoints.....	52
5.4 Launch the CGE Web Server .....	54
5.5 Create and Use a Database.....	54
6 Cray Graph Engine (CGE) Security Mechanisms.....	56
6.1 Grant Basic Access to Owned Query Engines .....	56
6.1.1 Eliminating the Interactive Host Key Verification.....	58
6.2 Create a CGE Specific RSA/DSA Host Key.....	59
6.3 Grant Other Users Access to Their CGE Query Engine .....	60
6.3.1 Grant Other Users Access to One of the Owned Data Sets .....	61
6.3.2 Grant Other Users Access to All of the Owned Data Sets .....	61
7 Built-in Graph Functions.....	63
7.1 Combining Graph Algorithms with SPARQL .....	63
7.2 Invocation of a Graph Function .....	64
7.2.1 The <code>CONSTRUCT</code> Clause .....	64
7.2.2 The <code>INVOKE</code> Clause.....	64
7.2.3 The <code>PRODUCING</code> Clause.....	65
7.3 Inputs to the Graph Function .....	66
7.4 Sequence of Operators .....	67
7.5 Bad Rank.....	68
7.6 Betweenness Centrality.....	69
7.7 Community Detection Label Propagation (LP).....	70

---

7.8 Page Rank.....	71
7.9 S-T (Source – Target) Connectivity.....	72
7.10 S-T Set Connectivity.....	73
7.11 Triangle Counting.....	74
7.12 Vertex Triangle Counting.....	75
7.13 Triangle Finding.....	76
8 Cray Graph Engine (CGE) Extension Functions.....	78
8.1 Cray Graph Engine (CGE) Interval Analytics Functions.....	78
8.2 Cray Graph Engine (CGE) Haversine Functions.....	82
8.3 Cray Graph Engine (CGE) Square Root Function.....	82
9 Cray Graph Engine (CGE) Property Path Support.....	84
10 Cray Graph Engine (CGE) Quick Reference.....	87
11 Use the Cray Graph Engine (CGE) for a Hello World Example.....	90
12 CGE Java API.....	93
12.1 Use CGE Java API via Maven.....	95
12.2 Use CGE API via Java Development Kit (JDK).....	97
12.3 Use CGE API via Pre-built Main Entry Points.....	98
12.4 Use Case: A Comprehensive Java Program.....	100
12.5 Limitations of CGE Java API.....	110
13 Logging and Troubleshooting.....	112
13.1 Troubleshooting Common Cray Graph Engine (CGE) Issues.....	113
13.2 Terminate Orphaned cge-server Jobs.....	119

# 1 About the Cray® Graph Engine User Guide

---

The Cray® Graph Engine User Guide contains information about using the Cray Graph Engine (CGE), its Command Line Interface (CLI) and Graphical User Interface (GUI) to create and use RDF databases.

## Release Information

This publication version addresses the product version 2.0.UP00 of the Cray® Graph Engine.

## Record of Revision

Publication Number	Date	Description
S-3010-115	March 2015	This publication addressed the Beta release of the Cray® Graph Engine.
S-3010-1000	March 2016	This publication addressed the 1.0.UP00 release of the Cray® Graph Engine.
S-3010-2000	August 2016	This publication addressed the 2.0.UP00 release of the Cray® Graph Engine.

## Typographic Conventions

Monospace	Monospaced text indicates program code, reserved words, library functions, command-line prompts, screen output, file names, path names, and other software constructs.
<b>Monospaced Bold</b>	<b>Bold monospaced</b> text indicates commands that must be entered on a command line or in response to an interactive prompt.
<i>Oblique or Italics</i>	<i>Oblique or italicized</i> text indicates user-supplied values in commands or syntax definitions.
<b>Proportional Bold</b>	<b>Proportional bold</b> text indicates a graphical user interface window or element.
\ (backslash)	A backslash at the end of a command line is the Linux® shell line continuation character; the shell parses lines joined by a backslash as though they were a single line. Do not type anything after the backslash or the continuation feature will not work correctly.
Alt-Ctrl-f	Monospaced hyphenated text typically indicates a keyboard combination.

## Scope and Audience

This publication does not include in-depth information about RDF and SPARQL. The intended audience of this publication is users and system administrators. It is assumed that all the commands documented in this guide are executed via the bash shell.

## 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.

## Feedback

Visit the Cray Publications Portal at <http://pubs.cray.com> and make comments online using the [Contact Us](#) button in the upper-right corner or Email [pubs@cray.com](mailto:pubs@cray.com). Your comments are important to us and we will respond within 24 hours.

## 2 About the Cray Graph Engine (CGE)

---

CGE is a highly optimized software application designed for high-speed processing of interconnected data. It features an advanced platform for searching very large, graph-oriented databases and querying for complex relationships between data items in the 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.

### 2.1 Cray Graph Engine (CGE) Features

CGE features include:

- Optimized query engine for high-speed parallel data analysis.
- Support for submitting queries, updates and creating checkpoints.
- A rich Command Line Interface (CLI).
- The CGE graphical user interface, which acts as a SPARQL 1.1 end point. This interface enables editing SPARQL queries or SPARUL updates and submitting them to the CGE database. It also accepts a set of commands that allow users to perform various tasks, such as creating a checkpoint on a database, setting Name Value Pairs (NVPs) to control certain aspects of data preprocessing, and query processing etc.
- SPARQL query language extension via the `INVOKE` and `PRODUCING` operators, which allow a classical graph algorithm to be passed an RDF graph and for the algorithm's results to be returned as data that is compatible with SPARQL 1.1. This enables graph algorithm library calls to be nested within a SPARQL query.
- Multi-user support.
- Compatibility with POSIX-compliant file systems.
- Database preprocessing to apply inference rules to the data, as well as to index the data.

### 2.2 Concepts of Operation

CGE's operational model is comprised of the following major components:

- The graph oriented database
- Resource Description Framework (RDF)

### 2.2.1 What the Cray Graph Engine (CGE) is Not: a Relational Database

Most modern database systems use a relational representation of their data. This means that data items are stored in tables, with each row of the table holding data items that are in some way related to each other. For example, all of the data items in the same row might be associated with the same person, as shown in the following table:

Employee ID	Given Name	Family Name	Date Hired	Job Position
29650	Georgia	Smith	11/17/2001	Eng5

In practice, a relational row like the one above is contained in a table of many similar rows:

Employee ID	Given Name	Family Name	Date Hired	Job position
29650	Georgia	Smith	11/17/2001	Eng5
10926	Alex	Jones	2/5/2008	Mktng3
72219	Paul	Anderson	8/21/2005	Admin2

....

One of these fields is called the "key" and is used as the basis for looking up data from any of the other fields. In this example, "Employee ID" would probably be used as the key. The column labels, "Employee ID", "Given name" etc. are implicit. They are not stored with the table, but with a database *schema* that is associated with the table. The schema defines each field in the relation.

The kind of information that may be associated with a scheme is shown below:

Field	Name	Datatype
0	Employee ID	Integer, min 0, max 99999
1	Given name	Character, String length < 30
2	Family name	Character, String length < 30
3	Date hired	Integer 1-12, Integer 1-31, Integer > 1985
4	Job position	Character, String length < 10

**NOTE:** The database schema shown above is used as an example and is entirely conceptual. There are typically many tables in a large relational database, each with its own defining schema.

### 2.2.2 What the Cray Graph Engine (CGE) is: a Graph-Oriented Database that Uses RDF

RDF is a data representation standard that allows data from different schemas to be merged. It accomplishes this by extending the linking structure of the Web using Uniform Resource Identifiers (URIs) in order to create triples to name a subject, an object, and the relationship or predicate between the two.

#### Resource Description Framework Triples

An RDF triple contains three components:



- the subject, which is an RDF URI reference or a blank node
- the predicate, which is an RDF URI reference
- the object, which is an RDF URI reference, a literal or a blank node

Hence, data items in RDF are always represented as a trio of character strings, referred to as the “*subject*”, “*predicate*” and “*object*” fields. Because they were originally intended to be unique across the Internet, components of RDF triples use the generic URI / IRI syntax (RFCs 3986 and 3987).

A triple holding the same kind of information shown in the previous relational example might look like the following:

```
<http://cray.com/example/employeeID#29650>      (subject)
<http://cray.com/example/hasGivenName>          (predicate)
"Georgia"^^<http://www.w3.org/2001/XMLSchema#string> (object)
```

**NOTE:** The above three statements should be entered on a single line and have been shown in separate lines in this document due to lack of space. Furthermore, the text: (subject), (predicate) and (object) in the above lines are shown in this document for clarity and are not part of an actual triple.

Note that the given name data item “Georgia” is expressed as an RDF literal: the value coupled with a URL-like string identifying its data type. RDF triples are intended to be self-identifying in two ways, both of which can be seen in this example:


1. The literal's data type is attached to it.
2. The predicate identifies the class of data that the object belongs to, information that in the case of relational data, is implicit in the schema and the data item's position in the tuple. For RDF triples, there is no schema. That type of identifying information is explicit, in the predicate of the triple.

As is illustrated below, any subject-predicate-object triple can also be viewed as a source vertex-edge-sink vertex component of a directed graph:

```
<http://...ID#29650> <http://.../hasGivenName> "Georgia"^^<http://
www....#string>
```

**NOTE:** The statements shown above should be entered on a single line and have been shown on separate lines due to lack of space.

Figure 1. RDF Triple Viewed as a Graph Component



CGE is designed to store and analyze datasets when the patterns of relationships and interconnections between data items are at least as important as the data items themselves. The SPARQL query language is convenient in the sense that it provides most of the same features as SQL for filtering, grouping, and updating database information. Unlike SQL, however, SPARQL also provides a powerful mechanism for specifying (in a query) a complex interconnection pattern to search for in the database. For indefinite-size patterns and aggregate

information that can not be expressed in SPARQL, CGE supports the capability of nesting a call to a classical graph analysis function within a SPARQL query.

Each subject-predicate-object relationship is an RDF triple. In CGE, each element in the internal representation of the database includes a graph field, which specifies the subset of the graph that the triple belongs to. If the graph field is left blank, the triple becomes part of the default graph. Typically this default, or unnamed, graph is the main data subset.

## 2.3 About SPARQL

SPARQL is an RDF query language developed for semantic database queries. SPARQL queries replace the table and schema format of relational SQL queries with RDF triples and ontologies, which define predicates and relationships.

Some SPARQL features are listed below:

- The ability to explore data by querying for unknown relationships.
- Implicit `JOIN` syntax, which reduces the overhead for processing a complex query with multiple `JOINS` to the equivalent of traversing a graph.
- The `GRAPH` keyword allows data to be queried along with its source, returning both the data that matches the query and the name of the graph that contains the data.
- Unlike the proprietary query languages used in many graph database systems, SPARQL is a standardized, non-proprietary query language.
- Tools and APIs for interacting with RDF/SPARQL systems are widely and freely available for all major programming languages and platforms

This release of the CGE software supports a subset of SPARQL 1.1. The following SPARQL 1.1 features are not implemented:

- The `SERVICE` keyword, for querying remote data.
- The `MD5`, `SHA1`, `SHA256`, `SHA384`, and `SHA512` encryption functions.
- The `UCASE` and `LCASE` functions, which return a string literal whose lexical form is the upper or lower case of the lexical form of the argument, are implemented for ASCII characters only.
- The property paths feature, which extends the predicate portion of the query, allowing more extensive search patterns without the overhead of additional `OPTIONAL` statements.

**NOTE:** Although CGE does not natively support the SPARQL 1.1 property paths feature, it does support certain types of property paths. CGE's property path support is currently experimental and should be used with care. Contact Cray Support for additional information.

To learn more about SPARQL, visit <http://jena.apache.org/tutorials/sparql.html>

## 2.4 System Architecture Overview

CGE is designed to provide performance and scalability on large, complex, interconnected databases. Its query engine is based on a data parallelism approach, in which the software strives to keep every processor busy on a roughly equal fraction of the data. The query engine is serviced by a user interface and a command line interface.

See [Cray Graph Engine User Interface](#) and [Cray Graph Engine \(CGE\) Command Line Interface](#) for more information.

CGE uses the open-source Jena ARQ SPARQL parser to parse each query or update, and its parser auxiliary software translates it into a lower-level representation that can drive the query engine. Query results are written to the file system in a tab-separated-values (.tsv) format. For convenience, a pointer to the results file is returned to the user when the query completes.

Extensive logging information is also written as the query or update progresses, as an aid to troubleshooting.

## 2.5 RDF and SPARQL Resources

Cray recommends the following resources for learning more about RDF and SPARQL:

### RDF Resources

- RDF primer at <http://www.w3.org/TR/rdf-primer>

### SPARQL Resources

- "SPARQL by Example", available at <http://www.cambridgesemantics.com>, is an excellent introductory tutorial written by Lee Feigenbaum of Cambridge Semantics and Eric Prud'hommeaux of W3C
- SPARQL Tutorial at <http://jena.apache.org>
- "Learning SPARQL", available at <http://www.learningsparql.com> by Bob DuCharme
- SPARQLer Query Validator at <http://sparql.org/query-validator.html>.
- SPARQL 1.1 query language tutorial at <http://www.w3.org/TR/sparql11-query>

### Semantic Web Resources

"Semantic Web for the Working Ontologist", available at <http://www.workingontologist.org> by Dean Allemang and James Hendler.

### 3 Building a Database

The Cray Graph Engine (CGE) is launched using the `cge-launch` command. When the CGE application is launched, a database directory is specified using the `-d` option of the `cge-launch` command. Initially, this directory contains RDF data in N-triples or N-quads format. When the application is first launched on a new database directory, the database is compiled and stored in an internal format in the same directory. Subsequent launches of the application using the `cge-launch` command with the same database directory will use the compiled database. The `update` command can then be used to add to or update an existing database. For more information about using the `cge-launch` and `update` commands, see [Cray Graph Engine \(CGE\) Command Line Interface](#) and the `cge-launch` and `update` man pages.

For CGE to recognize raw RDF data to be built, the data must be presented in one of three ways in this directory

1. in a single file called `dataset.nq` (for N-Quads form data)
2. in a single file called `dataset.nt` (for N-Triples form data)
3. in multiple files listed in a file called `graph.info`

#### Converting Data to RDF Triples

CGE reads RDF data in N-triples or N-quads format. There are many third party tools that may be used to convert data into RDF. [D2R](#) is often used to extract data from an RDBMS into RDF format. The [TopBraid Composer by TopQuadrant™](#) can also be used to convert Excel, TSV, UML, or XML data. Conversion of data to RDF is beyond the scope of this publication.

#### Building the Internal Representation

Once the data has been translated into RDF, the user must place the data in the directory where CGE will build its compiled database files. If the RDF is contained in a single file, the simplest method is to rename this file to `dataset.nt` or `dataset.nq`. A `dataset.nt` has NTriples format, whereas a `dataset.nq` file has NQuads format. On the other hand, if the RDF is found in more than one file, a file named `graph.info` will need to be created. This file contains a list of RDF files, one file per line. Each file name in `graph.info` may optionally be followed by a graph name. If a graph name is specified, the graph name is applied to any triples found in the corresponding RDF file.

Following is a sample of a `dataset.nt` file which has been extracted from the Lehigh University Benchmark (LUBM) synthetic dataset:

```
<http://www.Department14.University0.edu/GraduateStudent87>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#takesCourse>
<http://www.Department14.University0.edu/GraduateCourse17> .
<http://www.Department14.University0.edu/GraduateStudent87>
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#TeachingAssistant> .
<http://www.Department14.University0.edu/GraduateStudent87>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#teachingAssistantOf>
<http://www.Department14.University0.edu/Course6> .
<http://www.Department14.University0.edu/GraduateStudent87>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#takesCourse>
<http://www.Department14.University0.edu/GraduateCourse18> .
<http://www.Department14.University0.edu/GraduateStudent87>
```

```
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#GraduateStudent> .
<http://www.Department14.University0.edu/GraduateStudent87>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#name>
"GraduateStudent87" .
<http://www.Department14.University0.edu/GraduateStudent87>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#emailAddress>
"GraduateStudent87@Department14.University0.edu" .
<http://www.Department14.University0.edu/GraduateStudent87>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#undergraduateDegreeFrom>
<http://www.University843.edu <http://www.university843.edu/>> .
<http://www.Department14.University0.edu/GraduateStudent87>
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#advisor>
<http://www.Department14.University0.edu/AssistantProfessor6> .
```

**IMPORTANT:** Each predicate must appear on its own line. Some predicates are shown on multiple lines in the sample above due to lack of space.

The specification for NTriples can be found at <http://www.w3.org/TR/2004/REC-rdf-testcases-20040210/#ntriples>.

Following is a sample of a `graph.info` file:

```
# example graph.info file

# filenames can be absolute
/lustre/scratch/users/jdoe/database1/dbtriples1.nt

# or they can be relative to the database directory, which is where the graph.info file resides
database2/dbtriples2.nt

# they can specify a named subgraph with a URI
/lustre/scratch/users/jdoe/database3/dbquads3.nq <http://cray.com/namedGraphs/Graph3>
```

**NOTE:** Triples and quads are supported in both the `.nt` and `.nq` files. Quads in the RDF file are not affected by the optional graph name specified in the `graph.info` file. Lines containing only white space or lines beginning with the comment character ('#') are ignored. If the file is a mix of triples and quads, the triples become part of the graph specified in the `graph.info` file.

As mentioned earlier, when the application is launched via the `cge-launch` command. The `-d` parameter specifies the database directory.



**WARNING:** The `-d` parameter is mandatory. Launching CGE without specifying it produces an error.

This directory must already exist if it has been populated with `dataset.nt`, `dataset.nq`, rules and/or a `graph.info` file. If a compiled database is not present, a database is built using the `graph.info`, `dataset.nt`, or `dataset.nq` file in that directory.

When the database has been built, the following files are saved in the database directory:

- `dbQuads`
- `string_table_chars`
- `string_table_chars.index`
- `graph.info` file is created (if not already present), which is only used to load in a database from RDF files and is not used once the database is compiled.

Once the database has been built, CGE can begin executing queries and updates. When the application is subsequently launched via the `cge-launch` command specifying the same directory, the `dbQuads` file is detected, and the compiled database is read rather than the RDF.

**IMPORTANT:** Executing multiple update commands at a time is not supported currently. Updates should be split into separate files and/or submissions.

When loading a dataset, CGE will look for a dataset in the following places:

- If `dbQuads` exists, it will be used.
- If `dbQuads` does not exist, but `graph.info` exists, `graph.info` will be opened and read to obtain a list of source data files, which will then be used to build a new dataset.
- If neither `dbQuads` nor `graph.info` exist, but `dataset.nt` (or `dataset.nq`) exist, `dataset.nt` or `dataset.nq` will be used to build a new dataset.
- If none of the above files exist, CGE will fail.

In each of these cases, if the file exists but is in some way invalid, CGE will fail.

## Memory Requirements

- **Memory Requirement for reading a database from RDF** - The amount of memory required to read a database from RDF depends on the number of triples/quads in the database, the number of unique strings in the dictionary, and the length of those strings. As a rule of thumb, however, the main memory should be 4 times the size of the RDF file(s). For example, for a 100 GiB triples file, at least 400 GiB ( $4 * 100$ ) should be used.
- **Memory Requirement for loading a compiled database** - A compiled database consists primarily of the `dbQuads` files, containing the compiled quads, and the `string_table_chars` files, containing the dictionary. To enable CGE to load the database and execute meaningful queries, the main memory should be 20 times the sum of the sizes of `dbQuads` and the `string_table_chars` file. For example, if `dbQuads` is 32 GiB and `string_table_chars` is 256 GiB, at least  $(20 * (32 + 256))$  GiB of memory should be used.

## 3.1 About Rules Files

One way to greatly increase the knowledge contained in the database is to provide a set of inferencing rules. These rules are used during the database builds and in subsequent data updates (whether by SPARQL updates or by editing the database) to create new relationships between objects. Providing inferencing rules grants SPARQL queries access to inferred data, in addition to the raw data that was imported into the system.

### Forward vs. Backward Chaining

There are two types of chaining:

- **Forward Chaining** - In forward chaining, the inferencing rules are recursively applied to the database, creating new quads and adding them to the database. If *a* implies *b* and *a* is in the database, we add *b* to the database.
- **Backward Chaining** - Rather than pre-computing quads in the database as in forward chaining, with backward chaining the queries are modified to support those rules. If *a* implies *b* and a query searches for *b*, it is changed to search for (*a* UNION *b*).

CGE's rules inference engine does not implement backward chaining, but it implements a highly parallel form of forward chaining.

## 3.2 Create a Set of Inferencing Rules

Inferencing can be performed to generate additional relationships once the Cray Graph Engine (CGE) builds a database. CGE accomplishes this with a user defined rules file, which contains a set of rules specific to the data being processed. The rules file format and semantics are based on the Jena rules, documented at <http://jena.apache.org/documentation/inference/index.html#rules>. In this version of CGE there are certain limitations to these rules, which are described in [Limitations to Jena Rules Syntax](#).

The rules file has the form: one or more prefixes, followed by one or more rules

left-hand side quad(s) -> right-hand side quad(s)

Comments are denoted by a # character at the beginning of a line.

The quad, or quads, on the left-hand side of the -> are the quads that the inferencer will attempt to match to infer the quad, or quads, on the right-hand side of the ->. All of the left-hand-side rules must be satisfied in order for the inference to be made. Each rule must end with a period (.) and a newline character, and each rule must be on its own line. The inferencer does not recognize the escape character (\).

A quad takes the form:

```
(subject predicate object [graph])
```

It is mandatory to specify the subject, predicate and object. The graph field is optional. If a graph is not specified, the inferencer will use the default graph and the rule will apply only to triples in that graph. The subject, predicate and object fields can be any valid form of these fields as specified by the N-Quads grammar, except as described in [Limitations to Jena Rules Syntax](#) in CGE. The graph field in a quad has the same valid forms as an object. If a rule contains a URI, that URI must have existed in at least one of the data files that were included in the database. Alternatively, to apply a new ontology that was not in the original data files, create a new file that contains any new objects and predicates, and add that file to the database. The fields of a quad in a rule can also be variables, or shorthand versions of strings built from a specified prefix. A variable must begin with a ? character, followed by a valid name. A name can contain any of the following characters:

```
name := [a-zA-Z][_a-zA-Z0-9]*
```

To specify one or more prefixes at the beginning of a rules file, before any rules, use the following syntax:

```
@prefix prefix_name: <http://urlstring#>
```

A rules file does not have to use prefixes. However they can be used to simplify quads within rules. For example, prefixes are useful for creating shorthand versions of URIs that will be used repeatedly in the rules statements.

As with rules, each prefix must end with a period (.) and a newline, and each prefix must be on its own line.

The following prefix and rule examples are from the rule set used for the LUBM data.

### Inferencing a Database

When a database is built with inferencing enabled and a `rules.txt` file is found in the database directory, CGE will start applying the forward chaining rules found in that file to the triples/quads read from the RDF. The inferred quads are added to the in-memory database and stored in the compiled `dbQuads` file. If inferencing is enabled, the `rules.txt` file is also used when updating a database using SPARUL commands. As with any other quads added by the SPARUL commands, the inferred quads are added to the in-memory database but are not written to disk until the database is check-pointed.

For more information, see [About Rules Files](#)

**NOTE:** Inferencing is enabled by default and may be disabled by setting the value of the `cge.server.InferOnUpdate` control parameter to 0. Control parameters are configuration keywords that allow controlling server configuration settings. For more information, see [Control Options](#).

## Examples

### A prefix statement

```
@prefix ub: <http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
(?x rdf:type ub:Course) -> (?x rdf:type ub:Work) .
```

In this example the term `rdf:type` is shorthand for:

```
<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>.
```

The inferencer expands the prefixed version of the string to the full string when creating the rules used during inferencing. The rule in this example says that for a given triple `?x rdf:type ub:Course` in the default graph, infer a new triple `?x is-type ub:Work` and add it to the default graph, as shown in the next example.

### Inferring a new triple

Applying this rule:

```
(?x rdf:type ub:Course) -> (?x rdf:type ub:Work) .
```

to this triple in the data input:

```
<http://www.Department10.University0.edu/Course6> <http://www.w3.org/1999/02/22-rdf-syntax-ns#type> \
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#Course>
```

infers (and adds) this new triple to the default graph:

```
<http://www.Department10.University0.edu/Course6> <http://www.w3.org/1999/02/22-rdf-syntax-ns#type> \
<http://www.lehigh.edu/~zhp2/2004/0401/univ-bench.owl#Work>
```

### A rule to establish a hierarchy of types

The following rule shows one way that ontology rules are used to establish a hierarchy of data types.

```
(?x rdf:type ub:Faculty) -> (?x rdf:type ub:Employee) .
(?x rdf:type ub:Employee) -> (?x rdf:type ub:Person) .
```

A Faculty member is also an Employee, an Employee is also a Person, and so on. Such a rule eliminates the need to explicitly including each desired type for each such item in the database. Note that this rule did not use the graph field.

The following rule uses a variable for the graph field. This rule is excerpted from the RDFS rules file, which is based on some of the Jena rules for RDFS and OWL. The complete rules file is reproduced in [Sample RDFS Rules File](#) .

```
(?x ?a ?y ?g) (?a owl:inverseOf ?b ?g) -> (?y ?b ?x ?g) .
```

This rule is also an example of another way rules are used to establish relationships between triples in the database. This rule states that if two predicates `A` and `B` are defined to be inverses of each other and then if the triple `(X A Y)` appears in the database, then the system can infer that the triple `(Y B X)` is also there, or should be there.



## Cross-database rules

Another use of a rules file is to establish a relationship between triples in two different databases. For example, if one were extending a U.S.-based database with some additional data from France, it might streamline the process to include such rules as:

```
(<x.cray.eg.france#personne> <x.cray.eg.france#nom> ?name <x.cray.eg.frenchdb>) -> \
(<x.cray.eg.us#person> <x.cray.eg.us#name> ?name <x.cray.eg.usdb>) .
```

By this rule the fields in the quads are translated into their English counterparts, consistent with the data that is already in the American based database.

## 3.3 Sample RDFS Rules File

The following sample rules file is based on the Jena rules for RDFS and OWL. It is reproduced here courtesy of w3.org.

```
# These rules are based on the Jena rules for rdfs, plus some Jena rules
# for OWL.
#Line breaks inserted into some of these rules for formatting purposes.
#This was done for readability within this document, but is not valid syntax.

# Make a prefix for rdf:type. The IRI is defined by the SPARQL to be

# http://www.w3.org/1999/02/22-rdf-syntax-ns#type, which we can
# shorthand with rdf:type by defining a prefix for rdf:
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .

# Shorthand for rdfs

@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .

# Shorthand for owl

@prefix owl: <http://www.w3.org/2002/07/owl#> .

# Skip this one.

# [rdf1and4: (?x ?p ?y) -> (?p rdf:type rdf:Property), (?x rdf:type
rdfs:Resource), (?y rdf:type rdfs:Resource)]
# Add rule for rdfs 2:

# [rdfs2: (?x ?p ?y), (?p rdfs:domain ?c) -> (?x rdf:type ?c)]
(?x ?p ?y ?g) (?p rdfs:domain ?c ?g) -> (?x rdf:type ?c ?g) .

# [rdfs2a: (?x rdfs:domain ?y), (?y rdfs:subClassOf ?z) -> (?x rdfs:domain ?z)]
(?y rdfs:subClassOf ?z ?g) (?x rdfs:domain ?y ?g) -
> (?x rdfs:domain ?z ?g) .

# Add rule for rdfs 3:

# [rdfs3: (?x ?p ?y), (?p rdfs:range ?c) -> (?y rdf:type ?c)]
(?x ?p ?y ?g) (?p rdfs:range ?c ?g) -> (?y rdf:type ?c ?g) .

# [rdfs3a: (?x rdfs:range ?y), (?y rdfs:subClassOf ?z) -> (?x rdfs:range ?z)]
(?y rdfs:subClassOf ?z ?g) (?x rdfs:range ?y ?g) -> (?x rdfs:range ?z ?g) .
```

```

# Add rule for rdfs 5a:

# [rdfs5a: (?a rdfs:subPropertyOf ?b), (?b rdfs:subPropertyOf ?c) ->
#   (?a rdfs:subPropertyOf ?c)]
# (?a rdfs:subPropertyOf ?b ?g) (?b rdfs:subPropertyOf ?c ?g) -> (?a
# rdfs:subPropertyOf ?c ?g) .

# Add rule for rdfs 6:

# [rdfs6: (?a ?p ?b), (?p rdfs:subPropertyOf ?q) -> (?a ?q ?b)]
# (?a ?p ?b ?g) (?p rdfs:subPropertyOf ?q ?g) -> (?a ?q ?b ?g) .

# Skip this one.

# [rdfs7: (?a rdf:type rdfs:Class) -> (?a rdfs:subClassOf ?a)]

# Add rule for rdfs 8:

# [rdfs8: (?a rdfs:subClassOf ?b), (?b rdfs:subClassOf ?c) ->
#   (?a rdfs:subClassOf ?c)]
# (?a rdfs:subClassOf ?b ?g) (?b rdfs:subClassOf ?c ?g) -> (?a rdfs:subClassOf ?c ?
# g) .

# Add rule for rdfs 9:

# [rdfs9: (?x rdfs:subClassOf ?y), (?a rdf:type ?x) ->
#   (?a rdf:type ?y)]
# Put the quad with the most potential matches as the first quad to
# try and improve performance since since the first quads are handled
# in parallel.
# (?a rdf:type ?x ?g) (?x rdfs:subClassOf ?y ?g) -> (?a rdf:type ?y ?g) .

# Add rules for inverse property from owl.

# [inverseOf1: (?P owl:inverseOf ?Q) -> (?Q owl:inverseOf ?P) ]
# [inverseOf2: (?P owl:inverseOf ?Q), (?X ?P ?Y) -> (?Y ?Q ?X) ]
# We again process the quad that most likely will have the largest number
# of potential matches first (make it first quad in rule) to prevent
# potential performance problems.
# (?a owl:inverseOf ?b ?g) -> (?b owl:inverseOf
# ?a ?g) . (?x ?a ?y ?g) (?a owl:inverseOf ?b
# ?g) -> (?y ?b ?x ?g) .

# Add rule for owl transitive property.

# [transitiveProperty1: (?P rdf:type owl:TransitiveProperty),
#   (?A ?P ?B), (?B ?P ?C) -> (?A ?P ?C)]
# We again process the quad that most likely will have the largest number
# of potential matches first (make it first quad in rule) to prevent
# potential performance problems.
# (?a ?p ?b ?g) (?p rdf:type owl:TransitiveProperty ?g) (?b ?p ?c ?g) -> (?a ?p ?
# c ?g) .

# Skip this one.

# [rdfs10: (?x rdf:type rdfs:ContainerMembershipProperty) -> (?x
# rdfs:subPropertyOf rdfs:member)]

```

**NOTE:** Each prefix and rule must appear on its own line. Some prefixes and rules are shown on multiple lines in the sample above due to lack of space.

## 3.4 Limitations to Jena Rules Syntax

This release of CGE does not support all aspects of Jena syntax and semantics for rules. Specifically:

- The `@include` construct is not supported.
- Calls to functions or built-in primitives, such as `print`, `all`, or `max` are not supported.
- The `[...]` syntax is not supported, including named rules.
- Backward chaining is not supported. Furthermore, backward syntax (`<-`) cannot be used to express forward chaining.
- If multiple premises or conclusions (quads) are specified on either side of the `->` in a single rule, each pair must be separated by a space. The use of commas as separators is not supported.
- Native `UTF-8` is not supported in rules files, however Unicode characters are supported within URIs, where they are valid syntax.

**NOTE:** It is important to note that turning inferencing on/off is a database level setting. Turning inferencing on can negatively impact performance. When this setting is set to `true`, the inferencer will run during the first time that the database compiles and for subsequent updates. Since the whole database is examined when inferencing occurs, turning this feature on after a period of time during which it was turned off, will still affect the data that was loaded during the period when it was turned off. In other words, if a user turns inferencing off and then adds or updates data, that data will also be inferenced once the user turns the inferencing feature on again and performs another update. The current version of CGE supports the 2004 RDF NTriples standard. The 2014 RDF 1.1 NTriples standard is not supported by CGE.

## 4 Launch the CGE Server Using the `cge-launch` Command

---

The `cge-launch` command is used to launch the query engine and enables creating and building a database in a single step. This command handles the details of allocating batch resources (if necessary), setting up the launch environment, and composing a command line for the query engine on a given platform. This eliminates the need for a user to be aware of the details of a given platform's batch system or other considerations.

**NOTE:** Try relaunching CGE if the system displays an error message saying, "Server failed to start up" upon execution of the `cge-launch` command.

There are two pieces of information that need to be specified when starting up the query engine. These include:

- The path to the directory where the data set resides. This can be specified via the `-d` option of the `cge-launch` command.
- The path to a directory where the result files produced by queries need to be placed. This can be specified via the `-o` option of the `cge-launch` command.

Following is an example of using `cge-launch`:

```
$ cge-launch -o pathToResultsFile -d path -l logfile
```

In the preceding statement, *pathToResultsFile* is used as an example for the path to the directory that will contain the results of queries and/or updates, *path* is used as an example for the path to the database directory and *logfile* is used as an example for the log file that will contain the command and server output.

The `-o` and `-d` options of the `cge-launch` command accept:

- UNIX style pathnames as naming files on a POSIX compliant file system
- URLs of the following forms:
  - `file://unix_pathname` - This form is the equivalent of the Unix Style Pathname in URL form
  - `hdfs://name-server-address[:name-server-port-number]/HDFS_pathname` - This type of format indicates that a Hadoop Distributed File System (HDFS) file or directory is known to the specified name server and is located within that name-server's name space at *HDFS\_pathname*.

Both the aforementioned forms must refer to a file/directory that is shared across and equally accessible from all nodes. CGE will determine where to look for this file/directory based on recognizing one of the aforementioned path formats.

A few points to note while using checkpoints:

- If a full URL is used, the checkpoint is written exactly as specified by the URL, which means that an HDFS URL will cause the checkpoint to be written to the path specified in the URL on the HDFS file system described by the rest of the URL, and a FILE URL (i.e. `file:/path`) will be written to the POSIX filesystem at the pathname specified in the URL.
- If a relative path (i.e. a simple path with no leading `/` character) is used, the checkpoint will be written in a directory relative to the data directory used at CGE startup.

- If a full pathname but not a URL is specified, the pathname will be interpreted within the space specified by the URL of the data directory used at CGE startup, so, if CGE was started using an HDFS URL, the checkpoint will be written at the specified path within HDFS, if CGE was started with a simple pathname or FILE URL, the checkpoint will be written at the specified path within the POSIX file space.

The `--configFile` parameter specifies the pathname of the configuration file to be used by the server when setting up persistently configured settings. If this option is present, its value is used, otherwise the `$CGE_CONFIG_FILE_NAME` environment variable is taken from the invoking environment and used.

The `--configDir` parameter specifies the pathname of a directory containing configuration information. This information may include a `cge.properties` file, SSH keys and `authorized_keys` files and other CGE configuration. This directory, if specified, will be added to list of directories CGE searches for these files. The value of this parameter will be placed in the `$CGE_CONFIG_DIR_NAME` environment variable and passed to CGE. If this option is not specified and `$CGE_CONFIG_DIR_NAME` is set in the invoking user's environment, that value will be used instead. For more information, see the `cge-launch(1)` man page.

## 5 Mechanisms to Interact with the Cray Graph Engine (CGE) Database

The following mechanisms can be used to interact with the CGE database:

- CGE Graphical User Interface (GUI)
- CGE Command Line Interface (CLI)

### 5.1 Cray Graph Engine (CGE) Command Line Interface

The CGE CLI provides access to all the core functionality of the database via the command line. This interface is provided as part of the standard installation of CGE.

The list of available CGE CLI commands can be retrieved by executing the `cge-cli help` command without any options, as shown below:

```
$ cge-cli help
```

`cge-cli` commands are listed in the following table:

*Table 1. CGE CLI Commands*

Command	Description
<code>cge-cli checkpoint</code>	Requests creation of a checkpoint
<code>cge-cli echo</code>	Allows sending echo requests, which can be used to ping CGE to check if it is up and responding
<code>cge-cli fe</code>	Launches a web-based interface for accessing the server via a browser and provides SPARQL endpoints, which can be accessed via standard SPARQL APIs and tooling
<code>cge-cli help</code>	Displays help information
<code>get-configuration</code>	Determines the locations being searched for configuration files and the effective properties.
<code>cge-cli keyword-lookup</code>	Provides help with converting keywords between names and indexes so that the log options for using with other commands can be determined.
<code>cge-cli log-info</code>	Retrieves the current logging setup of the server

Command	Description
<code>cge-cli log-lookup</code>	Provides help with converting log levels between names and values so that the log options to use with other commands can be determined.
<code>cge-cli log-reconfigure</code>	Reconfigures the default logging setup of the server. The logging configuration changes persist until the server is shut down.
<code>cge-cli nvp-info</code>	Retrieves the current NVP setup of the server
<code>cge-cli nvp-reconfigure</code>	Reconfigures the default NVPs of the server. The NVP configuration changes persist until the server is shut down.
<code>cge-cli output-info</code>	Retrieves the current output directory for results from the server.
<code>cge-cli output-reconfigure</code>	Requests that the output directory for results be changed. The changes made persist until the server is shut down.
<code>cge-cli query</code>	Runs queries against the server, takes in SPARQL queries from files or from <code>stdin</code> only when no other query options are provided
<code>cge-cli sparql</code>	Runs a mixture of queries and/or updates against the server, takes in SPARQL queries/updates from files or from <code>stdin</code> only when no other input options are provided
<code>cge-cli update</code>	Runs updates against the server, takes in SPARQL updates from files or from <code>stdin</code> only when no other update options are provided
<code>cge-launch</code>	Launches the CGE Query Engine

Where more specific help for an individual command may be obtained by running the `cge-cli help` command, as shown in the following example:

```
$ cge-cli help checkpoint
```

### 5.1.1 Cray Graph Engine (CGE) Command Output

CGE CLI commands produce the following types of output:

- **Logging** - Provides diagnostic information about what a command is doing and is useful primarily for diagnosing any issues that may occur. All logging output goes to standard error.
- **Command Output** - Provides actual informational output of the command's status, such as query results, update success/failure etc. All command output is transmitted to the standard output.

As each type of output goes to a different output stream, output can easily be separated using standard shell redirection e.g.

```
$ cge-cli query example.rq > results.txt 2> query-client.log
```

The above example redirects the command output to the `results.txt` file and the logging to `query-client.log` file.

### 5.1.2 Cray Graph Engine (CGE) CLI Common Options

Certain options that are common to all commands and are provided by the CGE CLI are described in the following table:

*Table 2. Common Command Line Options*

Option	Argument(s)	Default Value	Example	Purpose
<b>Communication Options</b>				
<code>--db-host</code> <code>--dbhost</code>	Host	Localhost	<code>--db-host machine</code>	Specifies the host on which the database is running
<code>--db-port</code> <code>--dbport</code>	Port	3750	<code>--db-port 12345</code>	Specifies the port on which the database is running
<code>--i</code> <code>--identity</code>	Identity directory	<code>~/.ssh</code>	<code>-i /my/custom/identity</code>	Specifies the path to a SSH identity directory to use for authenticating to the server. When omitted, several default locations are tried and the first valid location is used
<code>--trust-keys</code>	N/A	N/A	<code>--trust-keys</code>	When this option is set, new host keys will automatically be trusted even when running in non-interactive mode. This is useful in environments where the database port (and thus the host and port combination required to trust the key for) may frequently change. This option should only be used when connecting to



Option	Argument(s)	Default Value	Example	Purpose
				trusted database servers.
<b>Client Logging Options</b>				
--debug --verbose	N/A	N/A	--verbose	Enables verbose mode, which includes setting the log level to <code>debug</code> . All logging output goes to <code>stderr</code> , allowing it to be separated from command output, which goes to <code>stdout</code> .  If the <code>--quiet</code> option is also specified, then the verbose mode takes precedence.
--quiet	N/A	N/A	--quiet	Enables quiet mode, which sets the log level to <code>error</code> , causes little/no logging to go <code>stderr</code> . All logging output is transmitted to <code>stderr</code> , allowing it to be separated from command output, which is transmitted to <code>stdout</code> .  If one of the verbose mode options is also specified, precedence is given to the verbose mode.
--trace	N/A	N/A	--trace	Enables trace mode, which includes setting the log level to <code>trace</code> . All logging output is transmitted to

Option	Argument(s)	Default Value	Example	Purpose
				<p><code>stderr</code>, allowing it to be separated from the command output, which is sent to <code>stdout</code>.</p> <p>If the <code>--quiet</code> option is also specified, precedence is given to the verbose mode</p>
<b>Server Configuration Options</b>				
<code>--nvp</code>	Name and value	N/A	<code>--nvp cge.DoMemoryLeakDetection 1</code>	Sets a NVP to send to the server as part of the request. Usually necessary only if asked by Cray support to enable advanced options for debugging an issue.
<code>--log-disable</code>	N/A	N/A	<code>--log-disable</code>	Disables all server side logging for the request
<code>--log-level</code>	<i>Log_level</i>	N/A	<code>--log-level 16</code> Supported log levels include: <ul style="list-style-type: none"> <li>• 0=None</li> <li>• 1=Off</li> <li>• 2=Error</li> <li>• 4=Warn</li> <li>• 8=Info</li> <li>• 16=Debug</li> <li>• 32=Trace</li> </ul>	Changes the server logging level for the request.  Supported values may be obtained by using the <code>log-lookup</code> command.
<code>--log-string</code>	<i>Log_string</i>	N/A	<code>--log-level Foo</code>	Specifies a string that will be included in every server log line pertaining to the request. This is useful if it is required to isolate and extract the log lines specific to a request.

Option	Argument(s)	Default Value	Example	Purpose
<code>--log-keyword-level</code>	<i>Keyword_level</i>	N/A	<code>--log-keyword-level 41 32</code>	<p>Changes the server logging level for a specific logging keyword. The database server uses a keyword-based system that enables extracting log levels specific to certain parts of the request processing or changing the log level for a specific keyword.</p> <p>Supported values may be obtained by using the <code>log-lookup</code> and <code>keyword-lookup</code> commands.</p>
<code>--log-global-keyword</code>	<i>Keyword</i>	N/A	<code>--log-global-keyword 41</code>	Specifies that a given keyword should be included in all log lines.
<b>Miscellaneous Options</b>				
<code>-h command</code> <code>--help command</code>	N/A	N/A	<code>--help checkpoint</code>	Prints the help information for the command rather than running the command
<code>--batch</code> <code>--non-interactive</code>	N/A	N/A	<code>--non-interactive</code>	<p>When set, this option guarantees that the script will never prompt the user for input, i.e. it will never use <code>stdin</code>. This may cause some commands to fail if they would require any user input other than the provided options. This is useful when invoking the CLI in a non-interactive context.</p>

Option	Argument(s)	Default Value	Example	Purpose
<code>--configDir</code> <code>--config-dir</code>	Directory	N/A	<code>--configDir /path/to/config</code>	Sets the first location to search for configuration files
<code>--</code>	N/A	N/A	<code>--</code>	Used to separate the options from the arguments to the command. This is useful if arguments may be mistaken for options. Any arguments seen after the <code>--</code> are treated as arguments even if they could otherwise be considered as options.

### 5.1.3 SSH Identities

SSH is used to encrypt communications with the database and to verify that a user is authorized to access a database. An identity directory or directories to be used explicitly can be specified via the `-i` or `--identity` option. If this option is not used, the following defaults are used:

1. If the environment variable `$CGE_CONFIG_DIR_NAME` is defined, search it for keys
2. If there is a `.cge` directory under a user's home directory (as defined by the `$HOME` environment variable), search it for keys
3. If there is a `.ssh` directory under a user's home directory (as defined by the `$HOME` environment variable), search it for keys

Only keys from the first directory found to contain keys will be used. Enabling verbose mode displays log output, detailing which keys are being used.

### 5.1.4 CGE Hadoop HDFS Configuration

In some instances, the CGE command-line interface will need access to the HDFS configuration in order to access data results and configuration files that may be located there. To do so, the value of the environment variable `HADOOP_CONF_DIR` will be inspected and if this is a valid directory, the relevant configurations files from this directory will be used. Otherwise the default location `/etc/hadoop/conf` will be searched.

### Usage of URL paths for HDFS and Lustre

When checkpointing to Lustre, specify a full URL to the Lustre file. The pathname specified is interpreted relative to the scheme and authority of the data directory URL. So, to checkpoint to a different scheme, specify a URL that specifies the different scheme. To checkpoint to Lustre from HDFS, the path

```
file:/mnt/lustre/my/data/directory
```

will inform the `checkpoint` command where to put the data.

- If a full URL is used, the checkpoint is written exactly as specified by the URL, which means that an HDFS URL will cause the checkpoint to be written to the path specified in the URL on the HDFS file system described by the rest of the URL, and a FILE URL (i.e. `file:/path`) will be written to the POSIX filesystem at the pathname specified in the URL.
- If a relative path (i.e. a simple path with no leading `/` character) is used, the checkpoint will be written in a directory relative to the data directory used at CGE startup.
- If a full pathname but not a URL is specified, the pathname will be interpreted within the space specified by the URL of the data directory used at CGE startup, so, if CGE was started using an HDFS URL, the checkpoint will be written at the specified path within HDFS, if CGE was started with a simple pathname or FILE URL, the checkpoint will be written at the specified path within the POSIX file space.

### 5.1.5 Cray Graph Engine (CGE) Properties File

A `cge.properties` file may be used in order to specify some options without having to explicitly state them with every command invocation. This properties file may exist in any of the following locations:

1. If the `--configDir` option is set, use the properties file from that directory if it exists
2. If the environment variable `$CGE_CONFIG_FILE_NAME` is defined and that file exists, use that file
3. If the environment variable `$CGE_CONFIG_DIR_NAME` is defined and it contains a properties file, use that
4. If there is a properties file in the working directory from which the command line interface was launched, use that
5. If there is a `.cge` directory under your home directory (as defined by the `$HOME` environment variable) and it contains a properties file, use that

Only the first file found will be used. Enabling verbose mode displays output detailing exactly which properties file (if any) are used. Additionally, the `get-configuration` command can be used to view more detail, such as the locations being searched, which file is used, and the effective properties.

If present, values from this file are used unless these are specifically overridden using command line options.

Currently the following properties are supported:

*Table 3. CGE Property Files*

Property	Value	Equivalent Command Line Option	Description
<code>cge.cli.db.host</code>	<i>Host</i>	<code>--db-host</code> <code>--dbhost</code>	Host name of a CGE server that the CLI will connect to if the <code>--dbhost</code> option is not used.
<code>cge.cli.db.port</code>	<i>Port</i>	<code>--db-port</code> <code>--dbport</code>	Port number of a CGE server that the CLI will connect to if the <code>--dbport</code> option is not used.

Property	Value	Equivalent Command Line Option	Description
cge.cli.trust-keys	True / False	--trust-keys	Eliminates the need for a first-time interactive CLI command each time you start using a server on a new TCP/IP port number combination.

Note that if there is a properties file that overrides the default value, it will be noted in the logging and will contain a warning to alert the user of the fact that they have set it in the properties file. Leaving an out of date properties file around can interfere with correct communications with the database server with no clear reason.

## Defining Command Aliases

The properties file may also be used to define command aliases. These are essentially shortcuts to other commands. An alias is defined in the following manner:

```
$ cge.cli.alias.algebra=compile -c algebra
```

This defines a new alias algebra which simply invokes the compile command passing in the `-c Algebra` option. The CLI can then be invoked using the following command:

```
$ cge-cli algebra example.rq
```

This would compile the given query into algebra and is equivalent to running the following command:

```
$ cge-cli compile -c algebra example.rq
```

Command aliases are subject to the following restrictions:

- Aliases cannot override built-in commands.
- Aliases cannot be defined recursively, which means that an alias cannot be defined in terms of another alias.

## Advanced Command Alias Definition

There are some more advanced functions that can be performed on aliases such as using positional parameters. For example, consider the following definition:

```
$ cge.cli.alias.c=compile -c $1
```

This creates the alias `c`, which again invokes the compile command. However, this time it uses a positional parameter for the value of the `-c` option. With this definition the CLI can be invoked in the following manner:

```
$ cge-cli c rpn example.rq
```

Here the first argument after the alias is injected into the expansion of the alias so this is equivalent to running the following:

```
$ cge-cli compile -c rpn example.rq
```

**NOTE:** If a positional parameter receives no value, it will be passed through as-is, which will likely result in parser errors.

### 5.1.6 Create Checkpoints Using the CGE `checkpoint` Command

The `checkpoint` command is used to request the creation of the checkpoint, which is a dump to disk of the current database state, optionally including a NQuads file that can be used to export the database to other tools. A checkpoint is a compiled database consisting of a `dbQuads`, `string_table_chars`, and `string_table_chars.index` file.

This command simply takes in a path to a directory in which to create the checkpoint. The checkpoint directory is given as a URI. This URI may be a full `file://` or `hdfs://` URL, or it may be a relative URI (e.g. a simple absolute or relative pathname) in which case it will be resolved relative to the base URI on the server (i.e. the current database directory).



**CAUTION:** If a relative path is used, the path will be evaluated relative to the data directory of the running CGE instance.

It is possible to checkpoint to the same data directory the user started from by using that directory's path as the path for the checkpoint. The checkpoint, if successful, will overwrite the existing `dbQuads`, `string_table_chars` and `string_table_chars.index` files, so that the next time you start from that directory you will get the new dataset.

Alternatively, it is also possible to checkpoint to another directory. If the directory already contains a dataset, and the checkpoint succeeds, the dataset will be overwritten.



**CAUTION:** If the data directory is being moved to a different location, make sure to shutdown any instance of CGE that was launched using that data directory before relaunching CGE.

### Examples

Following are a couple of examples of using the `checkpoint` command:

- Using a relative URL to a file

```
$ cge-cli checkpoint /lus/scratch/user/db/cp1
```

- Using a HDFS URL

```
$ cge-cli checkpoint hdfs:/user/db/cp1
```

- Using NQuads

If an NQuads file needs to be generated for use with other RDF and SPARQL tools, use the `-q` or `--quads` option of the `checkpoint` command, as shown in the following example:

```
$ cge-cli checkpoint --quads /lus/scratch/user/db/cp1
```

Once the checkpoint has been created, the system will display a message saying: `Checkpoint creation succeeded`

### Additional items to note while using checkpoints:

- If a full URL is used, the checkpoint is written exactly as specified by the URL, which means that an HDFS URL will cause the checkpoint to be written to the path specified in the URL on the HDFS file system described by the rest of the URL, and a FILE URL (i.e. `file:/path`) will be written to the POSIX filesystem at the pathname specified in the URL.

- If a relative path (i.e. a simple path with no leading / character) is used, the checkpoint will be written in a directory relative to the data directory used at CGE startup.
- If a full pathname, but not a URL is specified, the pathname will be interpreted within the space specified by the URL of the data directory used at CGE startup. Therefore, if CGE was started using an HDFS URL, the checkpoint will be written at the specified path within HDFS, otherwise if CGE was started with a simple pathname or FILE URL, the checkpoint will be written at the specified path within the POSIX file space.
- The `checkpoint` command allows overwriting existing checkpoints. However it will do so in such a way that it guarantees that this is an atomic operation. This means that either the checkpoint is overwritten and replaced, or the previous checkpoint will continue to exist.

For more information, see the `cge-cli-checkpoint(1)` man page.

### 5.1.7 Compile SPARQL Commands Using the CGE `compile` Command

The `compile` command is used to compile SPARQL commands into the logical and/or physical plans that the database server will use to execute the command. This can be useful for understanding how the system is interpreting and optimizing a query or update.

Following is an example of using the `compile` command:

```
$ cge-cli compile -c algebra example.rq
```

The preceding example would compile the SPARQL command found in the `example.rq` file into algebra form and display it to standard output. Multiple files can be specified in order to compile a large number of files at once.

### Compilation Modes

The `-c/--compiler-mode` option is used to specify the desired compilation output type. Supported values are as follows:

*Table 4. Compilation Modes*

Compilation Mode	Output Mode
<code>algebra</code>	The optimized SPARQL algebra for the query/update as text in SPARQL Set Expression ( <code>SSE</code> ) format. This can be thought of as the logical plan for the query.
<code>raw-algebra</code>	The unoptimized SPARQL algebra for the query/update as text in <code>SSE</code> format. This is the unoptimized logical plan for the query.
<code>rpn</code>	The physical plan for the query/update in binary form. Primarily intended for Cray developer use only.
<code>rpn-string</code>	The physical plan for the query/update in text. Primarily intended for Cray developer use only.
<code>all</code>	Produces all of the above.

This option may be specified multiple times to request multiple output formats. If the `all` option is also specified, it would supersede any individual format requests. The `-a` or `--all` options can also be specified as a shortcut for specifying the `-c all` option.



## Compilation Output

By default, compilation output goes to stdout and can be redirected to a file if desired. However, if multiple files need to be processed, or if more than one output type needs to be generated, then it is recommended to use the `-f` or `--files` option, which outputs a file for each input and compilation mode combination in the directory where the `cge-cli` command is being executed. The output file names are automatically generated based upon the input file name by replacing the extension with the appropriate extension for the output type:

Table 5. Compilation Output

Output Type	Output File Extension
algebra	.algebra
raw-algebra	.rawalgebra
rpn	.rpn
rpn-string	.rpnstring

For example, suppose that there is a file named `getTenRows.rq` that contains the following SPARQL query:

```
sparql query: select * {?s ?p ?o} limit 10
```

Now execute the `compile` command on the `getTenRows.rq`, as shown in the following example:

```
$ cge-cli compile -c all getTenRows.rq --files
0 [main] INFO com.cray.cge.parser.sparql.algebra.OpAsRpnMessage - Started Algebra to RPN message conversion
2 [main] INFO com.cray.cge.parser.sparql.algebra.OpAsRpnMessage - Finished Algebra to RPN message conversion (3 operations)
```

The above command would create the following four files:

- `getTenRows.rawalgebra`
- `getTenRows.rawalgebra`
- `getTenRows.rpn`
- `getTenRows.rpnstring`

For more information, see the `cge-cli-compile(1)` man page.

### 5.1.8 Check the Database State Using the CGE `echo` Command

The `echo` command is used to check that the database server is up and able to respond to requests.

As the name implies, the `echo` command simply sends some data to the database server and checks that the server echoes it back correctly.

An example of using the `echo` command is shown below:

```
$ cge-cli echo Test data
```

The above command sends the data `Test data` to the server. If the data is sent successfully, the system returns a message saying: `Echoed data received and validated successfully.`

## Generating Test Data

Use the `echo` command to generate some random data to send to the server. This can be used to test much larger requests than what would need to be manually typed in. For example:

```
$ cge-cli echo -g 8000
```

If the data is sent successfully, the system returns a messaging saying: "Echoed data received and validated successfully".

For more information, see the `cge-cli-echo(1)` man page.

### 5.1.9 Launch the CGE Web Server Using the `fe` Command

The `fe` command is used to launch a web server that provides a user interface and SPARQL endpoints to CGE.

In order to stream query results over HTTP, this command must be running on a host that has access to the same file system that the database server is writing results to. Typically, this means executing the `fe` command on a login node of the system running CGE. It should be noted that since it is often required to have the user interface available for a long period, it is recommended to launch it in the background so that it is resistant to terminal disconnects.

For example:

```
$ nohup cge-cli fe > web-server.log 2>&1 &
```

When the CGE user interface server has started, the system returns a message indicating that the server has started and is ready to accept HTTP requests.

Once the user interface has been launched, it is possible to access the SPARQL endpoints on the machine and port displayed in the log message.

## Verifying the Server Connection

It may be useful to verify that the database server is up and running when starting the web server. In this case, the `--ping` option can be used with the `fe` command, which makes the `fe` command check that the database server is up and running before launching the web server:

```
$ cge-cli fe --ping
```

## Running on a Different Port

Sometimes it may be necessary to run the web server on a different port. This can be achieved by using the `--server-port` option with the `fe` command (whose default value is 3756) to supply an alternative port on which to run the web server, as shown below:

```
$ cge-cli fe --server-port 12345
```

If an alternative port is chosen to run the web server, it is important to modify the URLs appropriately when accessing the user interface.

For more information, see the `cge-cli-fe(1)` man page.

### 5.1.10 Determine How Locations Are Being Searched Using the `get-configuration` Command

The `get-configuration` command is used to determine locations of CGE configuration files and the effective properties. This command does not communicate with the database. It inspects the user's local environment and provides information to help understand how configuration is being discovered. For example:

```
$ cge-cli get-configuration
```

The output of this command includes relevant environment variables and the location of found configuration files. The output also indicates whether or not the file was found. If the file was found, the path to that file is shown. Finally, all CGE related properties from that file are listed along with their values, as part of the output.

### 5.1.11 Display keyword ID and User Friendly Keyword Name Mappings Using the `keyword-lookup` Command

The `keyword-lookup` command provides the means to lookup mappings between keyword IDs and user-friendly keyword names. These can be used to find the values that need to be passed to the log options when invoking other commands. Unlike most of the other commands, the `keyword-lookup` command does not actually contact the database.

For example, use the following command to lookup a specific keyword ID:

```
$ cge-cli log-lookup 28
```

Alternatively, a keyword ID can be looked up based upon the keyword name using the `keyword-lookup` command, as shown in the following example:

```
$ cge-cli keyword-lookup QRY
```

The `keyword-lookup` command can be used without any arguments to display the full mapping of levels to names, as shown below:

```
$ cge-cli keyword-lookup
```

For more information, see the `cge-cli-keyword-lookup(1)` man page.

### 5.1.12 Retrieve Default Server Logging Information Using the `log-info` Command

The `log-info` command retrieves information about the server's default logging configuration.

**NOTE:** The information returned by the `log-info` command does not necessarily reflect the logging settings for individual requests since all commands may use the [Cray Graph Engine \(CGE\) Command Line Interface](#) to change the log configuration for specific requests.

An example of using the `log-info` command is shown below:

```
$ cge-cli log-info
```

Example output of the above command is shown below:

```
$ cge-cli log-info
0 [main] INFO com.cray.cge.cli.commands.AbstractSimpleCgeCommand - Making
```

```
request...
Server Log Configuration:
Version 1 - Printing Enabled - Default Level Info (8) - Keyword Levels Set {0-42}
```

In the above example, we can see that the server is configured with the default settings, as indicated by the text: Default Level Info (8). However, in other cases we might see different settings, as shown in the following example output:

```
$ cge-cli log-info
0 [main] INFO com.cray.cge.cli.commands.AbstractSimpleCgeCommand - Making
request...
Server Log Configuration:
Version 1 - Printing Enabled - Default Level Warn (4) - Keyword Levels Set {0-42}
Keyword TCP (Index 41) = Debug (16)
```

In the second example above, we can see that the default level has been turned down to Warn, but the TCP keyword is turned up to Debug. The server's default log configuration can be used via the `log-reconfigure` command if needed. For more information, see `cge-cli-log-info(1)` and `cge-cli-log-reconfigure(1)` man pages.

### 5.1.13 Lookup Mappings Between Log level Values and User Friendly Log Level Names Using the `log-lookup` Command

The `log-lookup` command provides the means to lookup mappings between log level values and user-friendly log level names. These can be used to find the values that need to be passed to the log options, when invoking other commands. It does so without contacting the database.

An example of using the `log-lookup` command for looking up the log level that has a value of 16 is shown below:

```
$ cge-cli log-lookup 16
```

An example of looking up a level based on the name follows:

```
$ cge-cli log-lookup Warn
```

Alternatively, use the `log-lookup` command without any arguments to retrieve the full mapping of levels to names, as shown in the following example:

```
$ cge-cli log-lookup
```

For more information, see the `cge-cli-log-lookup(1)` man page.

### 5.1.14 Change the Default Logging Configuration of the CGE Server Using the `log-reconfigure` Command

The `log-reconfigure` command changes the default logging configuration of the server.

The information returned by the `log-info` command does not necessarily reflect the logging settings for individual requests since all commands may use the [Cray Graph Engine \(CGE\) Command Line Interface](#) to change the log configuration for specific requests.

For example:

```
$ cge-cli log-reconfigure --log-level 16
```

The system will display a message if an incorrect value is specified for the log-level.

Upon successful execution of this command, the system returns the message: "Received success response".

**TIP:** It is recommended to verify that the log configuration changes have been implemented by using the `log-info` command. It may also be helpful to use the `log-lookup` and `keyword-lookup` commands to determine the values that need to be passed the options, in order to configure logging settings as desired.



**WARNING:** Do not set the server log levels to `DEBUG` or `TRACE`, especially, if the Cray Graph Engine (CGE) server is running with a large number of images.

For more information, see the `cge-cli-log-reconfigure(1)` man page.

### 5.1.15 Display Information About the Default NVP Configurations Using the CGE `nvp-info` Command

The `nvp-info` command retrieves information about the default NPV configuration of the server. The information retrieved by the `nvp-info` command does not necessarily reflect the NVP settings for individual requests since commands may change the NVP configuration for specific requests.

An example of using the `nvp-info` command is shown below:

```
$ cge-cli nvp-info
```

If the server's default NVP configuration needs to be changed, use the `nvp-reconfigure` command.

For more information, see the `cge-cli-nvp-info(1)` man page.

### 5.1.16 Change Default NVP Configurations Using the CGE `nvp-reconfigure` Command

The `nvp-reconfigure` command is used to change the server's default NVP configuration.

Upon successful execution of this command, the system returns a message saying: "Received success response". Configuration changes are not necessarily reflected in the NVP settings for individual requests since commands may change the NVP configuration for specific requests. It is recommended to use the `nvp-info` command to verify that the changes have taken effect, as shown below:

```
$ cge-cli nvp-info
```

Most of the supported NPVs have a defined range of acceptable values. Values specified outside of those ranges will be normalized into the range for that NVP. Unsupported NPVs are simply ignored, with a warning printed in the database logs and their values will not be stored by the server.

For more information, see the `cge-cli-nvp-reconfigure(1)` man page.

### 5.1.17 Display Server Output Directory Information Using the `output-info` Command

The `output-info` command retrieves information about the current output directory of the server. This is the directory to which the server writes query results to for later retrieval.

An example of using the `output-info` command is shown below:

```
$ cge-cli output-info
```

The `output-reconfigure` command can be used if it is required to change the server's output directory.

For more information, see the `cge-cli-output-info(1)` man page.

### 5.1.18 Change the Server's Output Directory Using the CGE `output-reconfigure` Command

The `output-reconfigure` command is used to change the servers output directory to which it writes query results for later retrieval. This directory is given as a URI, full `file://` or `hdfs://` URLs may be used. If a relative URI is given then this is resolved relative to the base URI of the server i.e., the current database directory.

An example of using the `output-reconfigure` is shown below:

```
$ cge-cli output-reconfigure /new/output/directory
```

**NOTE:**

After executing the `output-reconfigure` command, it is recommended to use the `output-info` command to verify that the changes have taken effect, as shown below:

```
$ cge-cli output-info
```

For more information, see the `cge-cli-output-reconfigure(1)` man page.

### 5.1.19 Execute Queries Using the CGE `query` Command

The `query` command is used to execute queries against the running database. This command can be used to execute a single query or a sequence of queries.

Queries that need to be executed may be specified in a number of ways:

- By providing a list of files, which contain lists of files containing queries to be executed
- By providing the names of query files directly
- Via `stdin` (only if no queries are specified in other ways and the `--non-interactive` option is not used)

The supported input methods have the precedence shown in the list above. This means that if any list files are specified, those queries are executed before any queries are specified directly. This command may only be used to execute SPARQL queries. To execute updates, use the `update` command or to execute mixtures of queries and updates use the `sparql` command.

An example of using the `query` command is shown below:

```
$ cge-cli query --list queries.txt extra-query.rq
```

The above command will execute all the queries specified in the `queries.txt` file before executing the query specified in the `extra-query.rq` file. Executing queries by default produces only information about where to obtain the results and not the result itself.

An example of using the `query` command is shown below:

```
$ cge-cli query types.rq
```

Here we can see that the database returns the following results information:

```
0 28 1756 0 file:///lus/scratch/rvesse/results/queryResults.2016-06-13T13.47.22Z000.28889.tsv
```

This is a simple tab separated string with the fields being as follows:

*Table 6. Status Descriptions*

Column Index	Information
0	Status - will be 0 for successful queries
1	Result count - number of results returned
2	Result size - results size in bytes
3	Execution time - query execution time in seconds
4	Results location - path to the file containing the results
5	Error message - should be blank for successful queries

## Results File Format

The file containing the results is in SPARQL Results TSV format and contains only the tabular results for the query. This means that if an `ASK/CONSTRUCT/DESCRIBE` query has been created, the results file will not contain the final results.

## Printing Results

This simple format makes it easy to process with standard command line tools,. For example, the following command can be used to show the results in the console:

```
$ cge-cli query --quiet types.rq | cut -d$'\t' -f 5 | xargs cat
```

As noted earlier, the results file contains only the tabular results for the query. If results of an `ASK/CONSTRUCT/DESCRIBE` query are desired to be printed, see the 'Streaming Results' section below.

## Streaming Results

As already seen, it is possible to use simple command line tools to extract and dump the query results to `stdout`. However, this only works for `SELECT` queries, and when the results can be accepted in SPARQL Results TSV format. If it is desired to retrieve the final results in an arbitrary format, the `--stream` option of the `query` command will need to be used. This option may only be used when executing a single query and it takes the MIME type of the desired results format.

```
$ cge-cli query --stream application/sparql-results+xml types.rq
```

Results are returned in SPARQL Results XML format. Supported formats include the following:

Table 7. Output Result Formats

Query Types	MIME Types	Output Format
ASK and SELECT	application/sparql-results+xml	SPARQL Results XML
	application/sparql-results+json	SPARQL Results JSON
	text/csv	SPARQL Results CSV
	text/tab-separated-values	SPARQL Results TSV
CONSTRUCT and DESCRIBE	application/n-triples	NTriples
	text/turtle	Turtle
	application/rdf+xml	RDF/XML
	application/rdf+json	RDF/JSON
	application/ld+json	JSON-LD

**NOTE:** Requesting a format that does not match the query type or is unknown will result in an error.

There are also three special values that may be passed to this option:

- text
- json
- xml

When these values are specified, the CLI will automatically select an appropriate text (line-based), JSON or XML output format in which to stream the results, while taking into account the type of query being evaluated. For example providing `--stream text` might produce SPARQL results TSV for an ASK/SELECT query but produce NTriples for a CONSTRUCT/DESCRIBE query. When these special values are used, the exact output format will not be known in advance but will be guaranteed to fall into the general format given.

## Making Multiple Queries

When multiple queries are executed, they are executed in the order specified (subject to the aforementioned precedence of list files over individual files) and the command will print a results header for each query.

```
$ cge-cli query types.rq list-graphs.rq ask-types.rq
```

It can be seen that a results header is retrieved for each query run.

For more information, see the `cge-cli-query(1)` man page

### 5.1.19.1 Common Errors Encountered Using the Cray Graph Engine (CGE) query Command

There are a number of common errors that may be encountered when using the `query` command:

- **Parser Errors** - Parser errors occur when a query is not in valid SPARQL 1.1 format. The error message details the unexpected token, along with what the parser was expecting at that stage of the parsing.
- **Unsupported Features** The system will also return an error message in cases where it is attempted to use an unsupported SPARQL 1.1 feature.



- **Communications Error** - Communication errors occur when the CLI cannot connect to the database. There are a variety of errors that can be encountered, depending on exactly what goes wrong with the communications between the CLI and the database server. If the same error is consistently returned, it is recommended to:
  - Add the `--trace` option to the command to get detailed information about the communications being attempted.
  - Review the log messages carefully.
  - Review the server logs if possible, as there are some situations which will manifest as client side communications errors that can be tracked down to the configuration of the server.

For more information, see [Troubleshooting Common Cray Graph Engine \(CGE\) Issues](#) on page 113

### 5.1.19.2 Cray Graph Engine (CGE) Optimizer Configuration

On rare occasions, it may be required to change the query optimizer configuration. This can be performed by using the `--opt-off` and `--opt-on` options. Both of these options take the name of an optimizer flag to disable/enable as desired.

The following example shows how to set the optimizer flag to `off`:

```
$ cge-cli query --opt-off optFilterPlacement types.rq
```

Executing the above statement would execute the query with the filter placement optimization disabled. If both the enabled and disabled flag options are specified, the flag will be considered as disabled. There are some flags whose values cannot be changed regardless of the options given.



**CAUTION:** Turning optimization off may result in significantly increased memory usage and/or performance degradation. Therefore, it is strongly recommended that the optimizer configuration be changed only when advised to do so by a Cray support engineer.

### 5.1.20 Shutdown the CGE Server Using the `shutdown` Command

The `shutdown` command is used to instruct the Cray Graph Engine (CGE) server instance to shut down.

An example of using the `shutdown` command is shown below:

```
$ cge-cli shutdown
```

The `shutdown` command requests that the server be shutdown gracefully. If this command is executed by the user that owns the server process, the user will receive a success message indicating that the server has shut down. On the other hand, if this command is executed by a user who does not own the server process, the system will return an error message.

**NOTE:** If the server is in a bad state, then this command will not succeed. Standard Linux techniques for killing an application process should be used in this case.

### 5.1.21 Execute Sparql Queries and Updates Using the `sparql` Command

The `sparql` command is used to execute queries and/or updates against the database. It can be used to execute a single query/update or to execute a whole sequence of queries and/or updates.

Queries and updates to be executed may be specified in a number of ways:

- By providing list files which contain lists of query and/or update files to be executed
- By providing the names of query and/or update files directly
- Via `stdin` (only if no queries/updates are specified in other ways and the `--non-interactive` option is not used)

The supported input methods have the precedence shown in the list above. This means that if any list files are specified, queries specified in those list files will be executed before any queries specified directly.

This command may be used to execute a combination of SPARQL queries and updates. Use the `query` command to execute SPARQL queries. Use the `update` command to execute SPARQL updates.

An example of using the `sparql` command is shown below:

```
$ cge-cli sparql --list commands.txt extra-command.ru
```

The above command will execute all the queries specified in the `commands.txt` file before executing the queries specified in the `extra-command.ru` file.

An example of using the `sparql` command is shown below:

```
$ cge-cli sparql list-graphs.rq create-graph.ru list-graphs.rq
```

Executing queries/updates using the `sparql` command produces the corresponding results for the command. This means that for queries it produces information about the results and for updates it produces a success/failure message as appropriate.

For more information about the `sparql` command, see the `cge-cli-sparql(1)` man page.

## Differences Between the `sparql` and `query` Commands

The major differences between the `sparql` and `query` commands include:

- The `sparql` command can run a mixture of queries and updates, whereas the `query` command can run queries only.
- The `query` command can stream results directly using the `--stream` option.

### 5.1.22 Execute Updates on a Database Using the CGE `update` Command

The `update` command is used to execute updates on a database. This command can be used to execute a single update or a sequence of updates. It supports all the common options described in [Cray Graph Engine \(CGE\) Command Line Interface](#).

Updates to be executed may be specified in a number of ways:

- By providing list files, which contain lists of update files to be run.
- By providing the names of update files directly
- Via `stdin` (only if no updates are specified in other ways and the `--non-interactive` option is not used)

The supported input methods have the precedence shown in the list above. This means that if any list files are specified, updates contained within those files will be executed before any updates specified directly.

This command may only be used to execute SPARQL updates. If it is required to execute queries, use the `query` command. To execute a combination of queries and updates, use the `sparql` command

An example of using the `update` command is shown below:

```
$ cge-cli update --list updates.txt extra-update.ru
```

The above statement will execute all the queries specified in `updates.txt` file before executing the query specified in the `extra-update.ru` file. Executing an update returns a message indicating whether the update succeeded or failed.

```
$ cge-cli update create-graph.ru
```

## Running Multiple Updates

If multiple updates need to be executed, they will be executed in the order specified (subject to the aforementioned precedence of list files over individual files) and the command will print a success or failure message for each update.

For example:

```
$ cge-cli update create-graph.ru drop-graph.ru
```

For more information, see the `cge-cli-update(1)` man page.

## 5.2 Access the Cray Graph Engine (CGE) Graphical User Interface

CGE provides a simple interface for access via a browser and also provides SPARQL 1.1 protocol compliant endpoints. The CGE user interface enables you to perform a number of tasks, including:

- Executing queries
- Executing updates
- Creating checkpoints on a database
- Using advanced options for viewing and editing server configurations, as well as for performing server NVP and logging configuration changes.

To access the CGE user interface, point the browser at: `http://machine:3756/dataset/`, where *machine* is the host name of the machine where the web server is hosted. Multiple instances of CGE can be launched on the same node at different ports.

**IMPORTANT:** The firewall configuration of the host machine must allow for port 3756 to be accessed externally or this will not work, unless the browser is running on the same host. If the site's firewall configuration does not permit this, SSH port forwarding can be used to forward the remote port to the local machine, as shown in the following example:

```
$ ssh machine -L 3756:hostname:3756
```

In the above example, *machine* is the machine running CGE's web server. The first 3756 is the local host port to connect to, whereas *hostname:3756* is the remote reference.

**NOTE:** The results format received in the browser is dictated by the HTTP Accept header that your browser sends (or conversely that your programmatic HTTP client sends). The 'Force text/plain as the response Content' option controls the Content-Type header that the front end responds with, which affects how the browser interprets the response. Depending on the browser if this option is disabled (the default) then this might mean that it downloads/offers to save the response to a file rather than displaying it in the browser, enabling the aforementioned

option changes the response Content-Type to always be text/plain regardless of what format the front end actually outputs which forces the browser to display the response in the browser itself. If it is needed to display the results in a different format, customise the HTTP Accept header accordingly, most browsers have some means to configure this. For example in Firefox navigate to **About>Config**. Click through the warning if it appears and then search for accept and edit the value of the `network.http.accept.default` setting to add the desired content types. The closest thing to plain text that the front end will produce is text/tab-separated-values. Most browsers include `application/xml` in their default accept header, which mean you will typically get SPARQL XML results by default (or RDF/XML if it were a `CONSTRUCT` query).

The CGE UI can then be accessed by pointing the browser at: `http://localhost:3756/dataset/`.

Upon successfully accessing the CGE user interface the following screen will be displayed:

Figure 2. Cray Graph Engine User Interface

The screenshot displays the Cray Graph Engine User Interface, which is divided into two main sections: the Query Interface and the Update Interface.

**Query Interface**

SPARQL Query:

[Large text area for entering a SPARQL query]

☐ Force text/plain as the response Content-Type (forces results to be displayed in browser)

**Server NVPs:**

[Text area for entering NVPs with example comments: # Enter NVPs one per line in properties file format e.g. #serv.DoMemoryLeakDetection=1, # Lines beginning with a # are comments]

**Server Logging Options:**

Server Log Level: Info

Server Log String: [Text area] (Printed on each server log line for this request)

☐ Disable all server logging for this request

**Update Interface**

SPARQL Update:

[Large text area for entering a SPARQL update]

Queries and/or updates can be executed directly from this page. Additional tasks can be performed using the links at the top of this interface.

## 5.2.1 Execute SPARQL Queries Using the CGE UI

### About this task

The Cray Graph Engine (CGE) **Query Interface** allows executing SPARQL queries on a loaded RDF database running within CGE.

Figure 3. Query Interface

The screenshot shows the 'Query Interface' page of the Cray Graph Engine. At the top, there is a navigation bar with links: 'Query Interface' (active), 'Update Interface', 'Checkpoint Interface', 'Server Information', and 'Edit Server Configuration'. Below the navigation bar, the page title is 'Query Interface'. Underneath, it says 'SPARQL Query:' followed by a large text area for entering the query. Below the text area, there is a checkbox labeled 'Force text/plain as the response Content-Type (forces results to be displayed in browser)'. Further down, there is a section titled 'Server NVPs:' with a text area containing example NVPs: '# Enter NVPs one per line in properties file format e.g.', '#cge.server.DoMemoryLeakDetection=1', '# Lines beginning with a # are comments', and '# You can also enter ARQ optimizer flags in the following form e.g.'. Below this is a section titled 'Server Logging Options:' with a 'Server Log Level:' dropdown set to 'Use Server Default' and a 'Server Log String:' text field. There is also a checkbox labeled 'Disable all server logging for this request' and a 'Run Query' button at the bottom.

### Procedure

1. To access CGE's **Query Interface**, point the browser at `http://machine:3756/dataset/query`, or select the **Query Interface** link from the list of links at the top of the UI.
2. To execute a SPARQL query, enter it in the **SPARQL Query** field. The check box under the **SPARQL Query** field can be selected to specify that the server should return the query results with a `Content-Type` header value of `text/plain`. This will force the browser to display the results in the browser, as many browsers will download the results rather than display them by default.
3. To execute a query, simply select the **Run Query** button, which will submit the query to the server and deliver the results to the browser. The user interface uses standard HTTP content negotiation to determine the format in which to return the query results. Most browsers receive results in an `XML/JSON` format.

## 5.2.2 Execute SPARQL Updates Using the CGE Update Interface

### About this task

The Cray Graph Engine (CGE) **Update Interface** allows you to execute SPARQL updates on a database. SPARQL update is a language extension to SPARQL 1.1 that makes it possible to make updates to an active RDF database, using SPARQL query syntax. Use the CGE **Update Interface** to perform a number of tasks, including updating the default database to add or remove RDF triples and quads, copying or moving the contents of one database to another, and performing multiple update operations in a single action.

### Procedure

1. To access CGE's **Update Interface**, point your browser at <http://machine:3756/dataset/update> or select the Update Interface link from the top on the main landing page of the CGE user interface.

Figure 4. CGE Update Interface

Query Interface | Update Interface | Checkpoint Interface | Server Information | Edit Server Configuration

### Update Interface

SPARQL Update:

Server NVPs:

# Enter NVPs one per line in properties file format e.g.  
#cge.server.DoMemoryLeakDetection=1  
# Lines beginning with a # are comments  
# You can also enter ARQ optimizer flags in the following form e.g.

Server Logging Options:

Server Log Level: Use Server Default ⓘ

Server Log String: (Printed on each server log line for this request)

☐ Disable all server logging for this request

Run Update

2. To execute a SPARQL update, enter the update statement into the **SPARQL Update** field.
3. Select the **Run Update** button to submit the update for processing. Once the system has finished executing the update, it will send either a success/failure message as appropriate.

### 5.2.3 Create a Checkpoint Using the CGE UI

#### About this task

When a database is started for the first time its initial state is considered to be a checkpoint. When a change is made to the database, its state can be preserved by creating a checkpoint. This preserves a copy of the previous in-memory database. Creating a checkpoint creates a persistent record of the database state, which is written to the database directory in a file named `export_dataset.nq`.

**NOTE:** Checkpoints can only be created on running databases. If there are any queries or updates executing, it is important to ensure that they finish executing before a checkpoint is created, otherwise the state of the database in the checkpoint may not contain the desired updates to it.

#### Procedure

1. To access the Cray Graph Engine (CGE) **Checkpoint Interface**, point the browser at `http://machine:3756/dataset/checkpoint`, where `machine` is the machine running CGE's web server. This brings up the **Checkpoint Interface**, as shown below:

Figure 5. Creating a Checkpoint

The screenshot shows the 'Checkpoint Interface' web form. At the top, there are navigation links: 'Query Interface', 'Update Interface', 'Checkpoint Interface' (which is highlighted), 'Server Information', and 'Edit Server Configuration'. The main title is 'Checkpoint Interface'. Below it, there is a 'Checkpoint Location:' label followed by a text input field. Underneath is a checkbox labeled 'Include a quads dump (NQuads) with the checkpoint?'. The next section is 'Server NVPs:', followed by a text area containing placeholder text: '# Enter NVPs one per line in properties file format e.g. #serv.DoMemoryLeakDetection=1 # Lines beginning with a # are comments #'. Below this is the 'Server Logging Options:' section, which includes a 'Server Log Level:' dropdown menu set to 'Info', a 'Server Log String:' text input field with a note '(Printed on each server log line for this request)', and a checkbox labeled 'Disable all server logging for this request'. At the bottom left of the form is a 'Create Checkpoint' button.

2. Specify a location for the checkpoint in the **Checkpoint Location** field. This is the directory where the checkpoint will be saved. The server will generate an error if this directory does not exist or is read-only.
3. Select the **Create Checkpoint** button to create the checkpoint. This will return a success/failure message as appropriate, as shown in the following example output:

```
Checkpoint created at /lus/scratch/cge/datasets/lubm/0/temp
```

## 5.2.4 Cray Graph Engine (CGE) Advanced Options

CGE provides a number of advanced options that can be used to change the behavior of the database server for a specific request. Some of these options impact the server, whereas others impact individual requests. The user interface for configuring advanced options is shown in the following figure:

Figure 6. Server Name Value Pairs

The screenshot shows a user interface for configuring advanced options. It is divided into two main sections: 'Server NVPs' and 'Server Logging Options'.

**Server NVPs:** This section contains a text area with the following text:
 

```
# Enter NVPs one per line in properties file format e.g.
#serv.DoMemoryLeakDetection=1
#
# Lines beginning with a # are comments
#
```

**Server Logging Options:** This section contains:
 

- A 'Server Log Level' dropdown menu currently set to 'Info'.
- A 'Server Log String' text input field, followed by the text '(Printed on each server log line for this request)'.
- A checkbox labeled 'Disable all server logging for this request'.

**NOTE:** Options provided in this section of the user interface are relevant only for the processing of the request under consideration and should be updated for each individual request. If it is desired to change the options for the database server as a whole, it will be required to use the interface described in the 'Edit Server Configurations Using the CGE UI ' topic of this publication.

### Server NVPs (Name Value Pairs)

In the **Server NVPs** section, NVPs can be specified to pass to the database server. These can be used to control behavior or enable additional debugging information.

**IMPORTANT:** In most cases, it will not be required to enter anything in the **Server NVPs** field, unless specifically instructed to do so by a Cray representative for gathering information to aid in diagnosing encountered issues.

### Server Logging Configuration

The **Server Logging Options** section provides options that allow configuring the amount of logging the database server will produce in the server side logs during the processing of a request. The desired logging level (i.e. log verbosity) can be selected from the **Server Log Level** drop down, which is followed by the **Server Log String field**, in which a log string can be entered. The log string can be up to 128 characters and will be included on each log line pertaining to the request. This is often useful for extracting all the log lines pertaining to a specific request.

Messages of types `INFO`, `WARNING`, and `ERROR` can be logged in the system, `INFO` being the default log level.

This interface also provides the option to disable logging for the request entirely, though it is generally recommended to avoid this option as it makes it difficult to monitor the status of the server while it processes queries.



## 5.2.5 View Server Configurations Using the CGE UI

### About this task

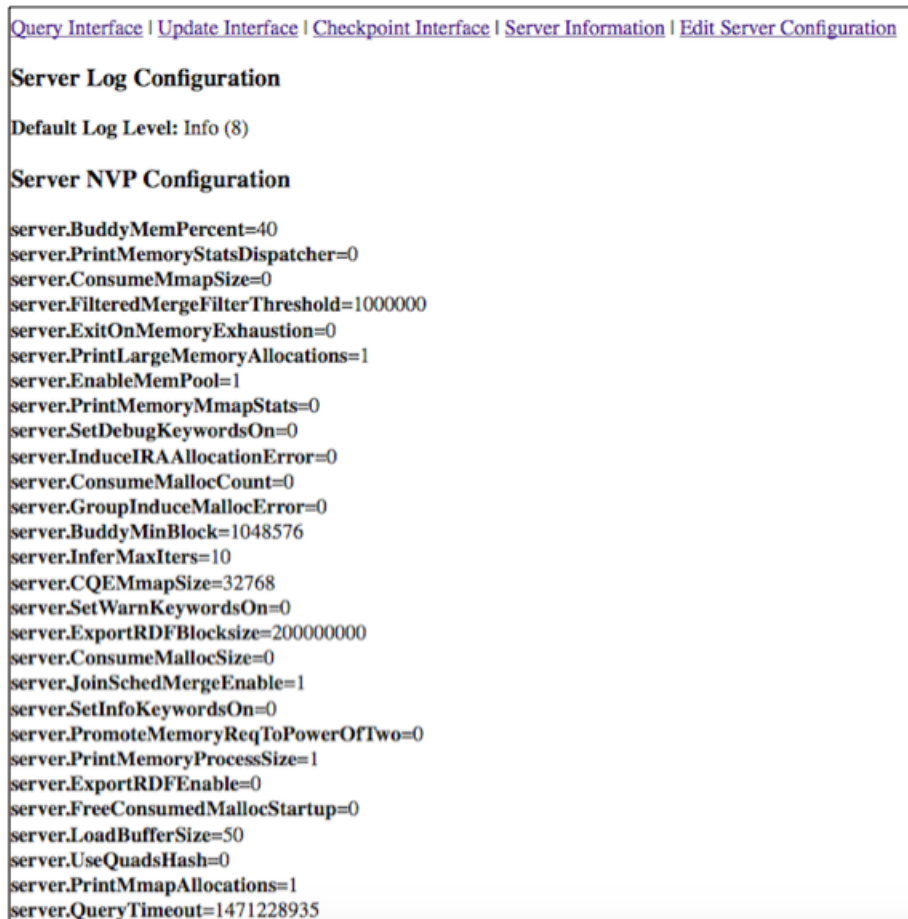
The **Server Information** interface allows you to view all the server configuration settings defined in the system.

**IMPORTANT:** In most cases it will not be required to change server configuration settings, unless specifically instructed to do so by a Cray representative in order to gather information for diagnosing issues being experienced.

### Procedure

To access the **Server Information** interface, point your browser at `http://machine:3756/dataset/info`, where *machine* is the machine running the Cray Gray Engine's (CGE's) web server. This displays a screen similar to the following:

Figure 7. Server Configurations



The screenshot shows a web interface with a navigation bar at the top containing links: [Query Interface](#), [Update Interface](#), [Checkpoint Interface](#), [Server Information](#) (which is highlighted), and [Edit Server Configuration](#). Below the navigation bar, the page is titled **Server Log Configuration** and shows **Default Log Level: Info (8)**. The next section is titled **Server NVP Configuration** and lists numerous server parameters and their values, such as `server.BuddyMemPercent=40`, `server.PrintMemoryStatsDispatcher=0`, `server.ConsumeMmapSize=0`, `server.FilteredMergeFilterThreshold=1000000`, `server.ExitOnMemoryExhaustion=0`, `server.PrintLargeMemoryAllocations=1`, `server.EnableMemPool=1`, `server.PrintMemoryMmapStats=0`, `server.SetDebugKeywordsOn=0`, `server.InduceIRAAllocationError=0`, `server.ConsumeMallocCount=0`, `server.GroupInduceMallocError=0`, `server.BuddyMinBlock=1048576`, `server.InferMaxIters=10`, `server.CQEMmapSize=32768`, `server.SetWarnKeywordsOn=0`, `server.ExportRDFBlockSize=200000000`, `server.ConsumeMallocSize=0`, `server.JoinSchedMergeEnable=1`, `server.SetInfoKeywordsOn=0`, `server.PromoteMemoryReqToPowerOfTwo=0`, `server.PrintMemoryProcessSize=1`, `server.ExportRDFFEnable=0`, `server.FreeConsumedMallocStartup=0`, `server.LoadBufferSize=50`, `server.UseQuadsHash=0`, `server.PrintMmapAllocations=1`, and `server.QueryTimeout=1471228935`.

The information displayed on the **Server Information** interface includes information about the log and NVP configurations of the server, as well as the results output directory.

## 5.2.6 Edit Server Configurations Using the CGE UI

### About this task

The **Edit Server Configuration** interface allows editing server configurations.

**IMPORTANT:** It is recommended not to change server configuration settings, unless specifically instructed to do so by a Cray representative in order to gather information for diagnosing issues.

### Procedure

1. To access the Cray Graph Engine (CGE) **Edit Server Configuration** interface, point the browser at `http://machine:3756/dataset/config`, where *machine* is the machine running CGE's web server. This presents the following screen:

Figure 8. Editing Server Configurations

2. Select the desired server NVP and logging options using the **Server NVPs** and **Server Logging Options** sections of the UI. In addition to the **Server NVPs** and **Server Logging Options**, this interface also contains a **Server Output Directory** field that allows changing the server output directory. This is the directory to which the database writes results, and from which the web server reads in order to deliver query results over HTTP.



**CAUTION:** It is best not to change server configuration settings, as doing so can adversely affect performance (especially if it is changed to point to a relatively slow file system).

3. When the required changes have been made, select the **Reconfigure Server** button to attempt the reconfiguration. Upon doing so, the system will return a response detailing the success/failures of the pieces of configuration that were to be updated, as shown in the following example output:

```
Updated Server NVP Configuration successfully
Updated Server Logging Configuration successfully
```

Unlike the options presented in the other interfaces, the values set from this interface persist for the lifetime of the server and become the new defaults.

### 5.2.6.1 Control Options

In most cases it will not be needed to change server configuration settings, unless a Cray support representative specifically requests, in order to gather information for diagnosing issues. However, there are some settings that you may occasionally wish to change. Name Value Pairs (NVPs) that enable you to modify these settings are listed in the following table:

Table 8. CGE NVPs

Parameter	Description	Default Value
<code>cge.server.QueryTimeout</code>	This parameter sets the maximum runtime (within the server) of a given query in seconds (wall clock time). This timeout will be checked after every operation. However, it does not interrupt operations. After the query times out, the server will terminate that query and will be immediately ready for subsequent queries. Acceptable values for this parameter range from 0 seconds (automatic termination at the start of the second operation) to 100,000 years expressed in seconds (3153600000000). If a negative value is entered for this field, it will be converted to 0.	31536000
<code>cge.server.InferOnUpdate</code>	Causes inferencing to be enabled or disabled for a given update. Has a value of either "0" or "1". The default value of this parameter is "1", which sets inferencing on for updates. A <code>rules.txt</code> file must be present for inferencing to take place. If no <code>rules.txt</code> file exists, inferencing will not be performed. If updates to the database were made after inferencing was turned on, triples added previously will stay saved in the database if inferencing is turned off subsequently.	1
<code>cge.server.BuddyMaxGBs</code>	Sets the upper limit on the amount of memory used by the big buddy allocator. The value of <code>BuddyMaxGBs</code> must be a non-negative integer value and is used to specify the maximum number of gigabytes allocated for the big buddy allocator. For example, setting the value to 50 will set the upper limit on the memory allocated for the big buddy allocator to 50 GB. By default, the limit is set to 128 GB and the maximum is 1 TB. Setting this parameter to 0 will disable the limit.	128 GB
<code>cge.server.LittleBuddyMaxGBs</code>	Sets the upper limit on the amount of memory used by the little buddy allocator. The value of <code>LittleBuddyMaxGBs</code> must be a non-negative integer value and is used to specify the maximum number of Gigabytes allocated for the little buddy allocator. For example, setting the value to 8 will set the upper limit on the memory allocated for the little buddy allocator to 8 GB. By default, the limit is set to 16 GB and the maximum is 128 GB. Setting this parameter to 0 will disable the limit.	16 GB

## 5.3 SPARQL Endpoints

Standard SPARQL tools can be used to interact with the Cray Graph Engine (CGE) by pointing them at the relevant endpoint URLs, which are shown in the following table:

Table 9. SPARQL Endpoints

Service	Endpoint URL
SPARQL Query	<code>http://machine:3756/dataset/query</code>
SPARQL Update	<code>http://machine:3756/dataset/update</code>

In the above examples, *machine* is used as an example for the name of the machine running CGE's web server. These endpoints are SPARQL 1.1 protocol compliant and provide all the standard parameters.

### Supported Content Types

The SPARQL query endpoint uses standard HTTP content negotiation to determine how to return query results to the SPARQL tool, depending on the **Accept** header that the tool sends.

**NOTE:** The results format received in the browser is dictated by the HTTP Accept header that your browser sends (or conversely that your programmatic HTTP client sends). The 'Force text/plain as the response Content' option controls the Content-Type header that the front end responds with, which affects how the browser interprets the response. Depending on the browser if this option is disabled (the default) then this might mean that it downloads/offers to save the response to a file rather than displaying it in the browser, enabling the aforementioned option changes the response Content-Type to always be text/plain regardless of what format the front end actually outputs which forces the browser to display the response in the browser itself. If it is needed to display the results in a different format, customise the HTTP Accept header accordingly, most browsers have some means to configure this. For example in Firefox navigate to **About>Config**. Click through the warning if it appears and then search for accept and edit the value of the `network.http.accept.default` setting to add the desired content types. The closest thing to plain text that the front end will produce is text/tab-separated-values. Most browsers include `application/xml` in their default accept header, which mean you will typically get SPARQL XML results by default (or RDF/XML if it were a `CONSTRUCT` query).

The following standard formats are supported by the query endpoint:

Table 10. Query Types and Supported Content Types

Query Type	Supported Content Types
ASK and SELECT	<ul style="list-style-type: none"> <li>SPARQL Results XML</li> <li>SPARQL Results JSON</li> <li>SPARQL Results CSV</li> <li>SPARQL Results TSV</li> </ul>
CONSTRUCT and DESCRIBE	<ul style="list-style-type: none"> <li>NTriples</li> </ul>

Query Type	Supported Content Types
	<ul style="list-style-type: none"> <li>• Turtle</li> <li>• RDF/XML</li> <li>• RDF/JSON</li> <li>• JSON-LD</li> </ul>

Standard HTTP behavior of returning the message "406 Not Acceptable" will apply if the tool does not include any formats the endpoint can produce in its **Accept** header.

## Custom Parameters

CGE features endpoints that provide custom parameters, which can be used to configure the same advanced options supported by the CGE user interfaces. These parameters are listed in the following table:

*Table 11. Custom Parameters*

Parameter	Example	Purpose
<code>forcePlainText</code>	<code>forcePlainText=true</code>	<p>Used to force the returned Content-Type to be text/plain regardless of the actual content type being returned.</p> <p>This is only useful for browser access to the endpoints and may cause errors if used with SPARQL tools.</p>
<code>nvps</code>	<code>nvps=foo%3Dbar</code>	<p>Specifies the NVPs to be passed to the database and applied to the request.</p> <p>These must be specified in Java properties file style with one <i>name=value</i> pair per line</p>
<code>log-level</code>	<code>log-level=16</code>	<p>Specifies the log level to use for database logging of the request. This takes an integer value with values interpreted as follows:</p> <ul style="list-style-type: none"> <li>• 2 = Error</li> <li>• 4 = Warn</li> <li>• 8 = Info</li> <li>• 16 = Debug</li> <li>• 32 = Trace</li> </ul> <p>The <code>log-lookup</code> command can be used for translating integer values to the desired log levels.</p>
<code>log-string</code>	<code>log-string=Foo</code>	Specifies a string to be included on every database log entry pertaining to the request.

Parameter	Example	Purpose
		Maximum supported length is 128 characters and longer strings will be truncated accordingly.
log-disable	log-disable=true	Can be set to disable all database logging for the request

## 5.4 Launch the CGE Web Server

Before using the Cray Graph Engine GUI, it is required to launch the database via the `cge-launch` command and leave the default port setting of 3750 unchanged. If an alternative port has been used, then it will be required to add the `--db-port` option to specify an alternative port. Once the database has been launched, the Cray Graph Engine (CGE) graphical user interface and/or the SPARQL endpoints may be used. This can be accomplished by launching the web server that provides the user interface on a login node of the system where CGE is running, as shown below:

```
$ cge-cli fe
```

Alternatively, you can use the following command to have the web server continue running in the background with its logs redirected, even if you disconnect from the terminal session:

```
$ nohup cge-cli fe > web-server.log 2>&1 &
```

**NOTE:** The web server is launched by the same script as the rest of the Command Line Interface tools, and supports many of the same standard options detailed in [Cray Graph Engine \(CGE\) Command Line Interface](#).

## 5.5 Create and Use a Database

### Prerequisites

If the Cray Graph Engine (CGE) is needed to perform inferencing on data, ensure that a valid `rules.txt` file exists in the directory containing the data.

### About this task

The following instructions can be used to create a database and execute queries and/or updates on the database once it has been built.

### Procedure

1. If the data is not in RDF format, convert the data to RDF.
2. If the RDF data resides in a single file, save/rename that file to `dataset.nt` or `dataset.nq`. This is required because CGE accepts ONLY files in `.nt` or `.nq` formats as input. All other formats should be

converted to either `.nt` or `.nq` (including `.rdf`). On the other hand, if the data resides in more than one file, create a `graph.info` file and add the names of the RDF file to that file.

3. Build the database using the `cge-launch` command as shown below:

```
$ cge-launch -o pathtoResultsDir -d path -l logfile
```

In the above statement, *pathtoResultsDir* is used as an example for the path to the directory that will contain the results of queries and/or updates. *path* is used as an example for the path to the database directory and *logfile* is used as an example for the log file that will contain the command and server output. *pathtoResultsDir* **MUST** be a directory and **MUST** contain either a triples or quads file. These files must be named `dataset.nt` or `dataset.nq` respectively. For more information, see the `cge-launch(1)` man page.

**NOTE:** When the database has been built, the following files are saved in the database directory:

- `dbQuads`
- `string_table_chars`
- `string_table_chars.index`

Executing multiple update commands at a time is not supported currently. Updates should be split into separate files and/or submissions. Collectively, the aforementioned files are the disk representation of the binary version of the database which can be reloaded into CGE. When the CGE application is launched again and the same database directory is specified, the `dbQuads` file will be detected and the compiled database will be read instead of the RDF. Furthermore, if the database directory contains a `rules.txt` file, CGE will perform inferencing on the data. This is because inferencing is turned on by default. It can be turned off by setting the `cge.server.InferOnUpdate` NVP parameter to 0.

4. Execute the `fe` command to launch a web server that provides a user interface and SPARQL endpoints to CGE.

```
$ nohup cge-cli fe > web-server.log 2>&1 &
```

5. To execute a query or update on the database, use either the CGE UI or the CGE CLI.
  - a. To execute queries/updates via the CGE UI, follow the instructions listed below:
    1. Connect to the CGE UI by pointing the browser at: `http://machine:3756/dataset/`. This brings up the CGE UI.
    2. Select the **Query Interface** or **Update Interface** to execute queries and updates respectively. Optionally, server configuration parameters can also be specified to control the query/update.
  - b. To execute queries/updates via the CGE CLI, use the `query`, `update` and `sparql` commands to execute SPARQL queries, updates and/or combination of queries and updates correspondingly. For usage information, see the associated man pages.

## 6 Cray Graph Engine (CGE) Security Mechanisms

---

The CGE query engine protects the port on which it communicates with clients using an encrypted authentication mechanism based on the Secure Shell (SSH) passwordless authentication mechanism. Before using the CGE user interface query clients to make requests on data sets, authentication must be configured. If it is required to set up the query engine to permit multiple users to execute requests, it will be required to configure public keys for each user. This can be configured on a per-data set or all data sets basis.

### 6.1 Grant Basic Access to Owned Query Engines

#### About this task

The Cray Graph Engine (CGE) query engine and CGE CLI commands use your SSH configuration to obtain public and private keys for use in authentication. Configuring basic query engine authentication is almost the same as configuring SSH passwordless authentication to the localhost IP host for your login account. The steps involved in granting basic access to your query engine are listed below:

#### Procedure

1. Ensure that you have a `.ssh` directory in your home directory and that the directory permissions are `700` (`rwX-----`).

To find out whether you have a `.ssh` directory, and whether or not it is correctly protected, use the following command:

```
$ ls -ld $HOME/.ssh
drwx----- 6 username group 204 Nov 20 07:15 /users/username/.ssh
```

If this looks correct you can move on to the next step. If the directory does not exist at all, you will need to create it, as shown below:

```
$ mkdir $HOME/.ssh
$ chmod 700 $HOME/.ssh
$ ls -ld $HOME/.ssh
drwx----- 6 username group 204 Nov 20 07:15 /users/username/.ssh
```

If the directory does not have the correct permissions, you can simply change those. However, it is important to ensure that the directory is writable only by you. As long as this requirement is met, you do not need to change anything. The following command can be used if it is required to set the permissions on the directory:

```
$ chmod 700 $HOME/.ssh
$ ls -ld $HOME/.ssh
drwx----- 6 username group 204 Nov 20 07:15 /users/username/.ssh
```



2. Create a public / private authentication key pair using `ssh-keygen` if the key pair does not currently exist. Use the following command to find out whether or not a public / private key pair has been configured.

**NOTE:** The following shows only key files (there will probably be other files as well unless this is a brand new `.ssh` directory):

```
$ ls -l $HOME/.ssh
total 80
-rw----- 1 username group 668 Apr 8 2014 id_dsa
-rw-r--r-- 1 username group 601 Apr 8 2014 id_dsa.pub
-rw----- 1 username group 883 Apr 8 2014 id_rsa
-rw-r--r-- 1 username group 221 Apr 8 2014 id_rsa.pub
```

In the above example, there may be only an RSA key pair (`id_rsa` and `id_rsa.pub`), only a DSA key pair (`id_dsa` and `id_dsa.pub`) or both. A file with ".pub" in its name is a public key file. A file without ".pub" in its name is a private key file. All of your private key files should have `-rw-----` for their permissions as shown above. Your public key files may be readable (not writable) by anyone, but do not need to be, so the permissions shown above are okay, but not required. The minimum permission set that should be used is `-rw-----`, this enables reading and modifying the file. The maximum permission set should have `-rw-r--r--`, which permits other users to read but not modify the public key. If there is not even a single public/private key pair in the `.ssh` directory, an SSH key will need to be generated. This can be done using the `ssh-keygen` command:

```
$ ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/users/username/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /users/username/.ssh/id_rsa.
Your public key has been saved in /users/username/.ssh/id_rsa.pub.
The key fingerprint is:
eb:0d:10:cd:4f:4b:f1:2b:20:87:99:82:93:b5:8d:ee [MD5] username@host
The key's randomart image is:
+--[ RSA 2048 ]-----+
|      .               |
|    + + *   o         |
|  + + B = o .         |
|   o . + = . .        |
|    . . S + .         |
|      . . .           |
|   E   o              |
|      . o             |
|      . .             |
+--[ MD5 ]-----+
$ ls -l $HOME/.ssh
total 8
-rw----- 1 username group 1679 Jan 6 11:49 id_rsa
-rw-r--r-- 1 username group 391 Jan 6 11:49 id_rsa.pub
```

This produces a public / private key pair which can be used for passwordless authentication to localhost.

**NOTE:** At present, CGE does not support `ssh-agent` forwarding, so it is not recommended to specify a pass-phrase when creating a key.

3. Place the public authentication key in the `.ssh/authorized_keys` file. This will enable interacting with CGE query engines started by the user on this machine (it does not allow other users to use the user's query engines). Set this up as follows:

```
$ cat $HOME/.ssh/id_*.pub >> $HOME/.ssh/authorized_keys
$ chmod 644 $HOME/.ssh/authorized_keys
$ ls -l $HOME/.ssh
total 80
-rw-r--r-- 1 username group 2601 Jun 18 2014 authorized_keys
-rw----- 1 username group 668 Apr 8 2014 id_dsa
-rw-r--r-- 1 username group 601 Apr 8 2014 id_dsa.pub
```

```
-rw----- 1 username group 883 Apr 8 2014 id_rsa
-rw-r--r-- 1 username group 221 Apr 8 2014 id_rsa.pub
```

4. Test using ssh to log into localhost without a password. The simplest way to test this is to try connecting to localhost through SSH. This will have the effect of logging on to the same host the the user is currently logged on to:

```
$ ssh localhost
The authenticity of host 'localhost (::1)' can't be established.
ECDSA key fingerprint is 0a:34:d6:d9:71:b4:6c:e6:1d:49:95:ea:7d:09:54:89 [MD5].
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'localhost' (ECDSA) to the list of known hosts.
Last login: Tue Jan 6 11:56:10 2015 from localhost
-----
Message of the day...
-----
$ exit
```

As you can see, the first time you do this, you will be prompted to verify that the key for localhost is correct. The user will also be prompted like this the first time the user tries to connect with a query engine with a new TCP/IP port number, so it is a good idea to do an interactive query or other kind of front-end command before trying to use a new query engine port from a script or other automated environment. Once authenticity of the host / port pair has been verified, this pair will be added automatically to your list of known hosts and the user should not need to do this again. To avoid the need for performing the interactive Host Key verification step, see [Eliminating the Interactive Host Key Verification](#) To show that this works, try a second attempt to use SSH to log into localhost:

```
$ ssh localhost
Last login: Tue Jan 6 11:56:10 2015 from localhost
-----
-
Message of the day...
-----
-
$ exit
$
```

5. Once this has been set up, it is required to authenticate the localhost / <port number> pairs for all query engine ports so that the clients can connect non-interactively. To do this, start CGE on each port you intend to use and run an interactive request through CGE, once for each port. The cge-cli echo command provides a simple way of doing so, as shown in the following example:

```
$ cge-cli echo --db-port=73737
The authenticity of host 'localhost' can't be established.
RSA key fingerprint is d2:b4:ad:70:f1:44:d3:8a:f5:16:db:db:76:07:19:47.
Are you sure you want to continue connecting? [Yes/No]: yes
13835 [main] WARN com.cray.cge.cli.communications.client.ssh.LoggingBridge - Permanently added 'localhost' (RSA) to the list of
known hosts.
14110 [main] INFO com.cray.cge.cli.commands.debug.EchoCommand - Sending echo request...
14157 [main] INFO com.cray.cge.cli.lightweight.commands.debug.EchoCommand - Echoed data received and validated successfully
```

To avoid the need for performing the interactive Host Key verification step, see [Eliminating the Interactive Host Key Verification](#)

### 6.1.1 Eliminating the Interactive Host Key Verification

The SSH protocol uses the host key to authenticate the server to the client, which is of particular importance when the client will be sending confidential data (passwords, for example) to the server. Since the SSH protocol used by CGE does not permit the use of passwords, and the clients do not generally send other secrets to CGE, there is no real need for the client (and the invoking user) to verify that the host key is the one that the user trusts.

By default, the CGE CLI commands require explicit first time verification of host keys, as you have seen in the examples above. There is, however, a setting that you can set in your cge.properties file(s) that will cause the

CGE CLI commands to consider any host key as trusted. This eliminates the need for a first-time interactive CLI command each time you start using a server on a new TCP/IP port number, and streamlines the process of connecting to a new instance CGE.

To add this setting, make sure that all appropriate `cge.properties` files contain the following line:

```
cge.cli.trust-keys=true
```

The same behavior can be achieved by adding the `--trust-keys` option to any of the CGE CLI commands.

**IMPORTANT:** While implicitly trusting host keys for CGE is generally a safe practice, in the case where your data set contains actual confidential data, and you are using the CGE CLI clients to update the data set with new confidential data, you want to be certain that there is nothing other than CGE itself listening to the contents of your updates. In that case, the host key is an important part of ensuring that there is nothing between you and your CGE instance. This is not expected to be a common case among CGE users, but if your use of CGE falls into this category, it is recommended not to use the mechanisms described here.

## 6.2 Create a CGE Specific RSA/DSA Host Key

### About this task

At some sites, site policy may dictate the use of a pass phrase with SSH keys used for logging into a system. If a pass phrase is used when creating your SSH key, the CGE authentication mechanism will be unable to use your SSH key(s) as its host key(s), so separate CGE specific host key(s) will need to be created. To do this, follow the instructions listed below:

### Procedure

Create the key in the `.cge` directory using `ssh-keygen (1)` instead of creating the key in the `.ssh` directory:

```
$ mkdir -p $HOME/.cge
$ ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/users/username/.ssh/id_rsa): /users/username/.cge/id_rsa
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /users/username/.cge/id_rsa.

Your public key has been saved in /users/username/.cge/id_rsa.pub.
The key fingerprint is:
eb:0d:10:cd:4f:4b:f1:2b:20:87:99:82:93:b5:8d:ee [MD5] username@host
The key's randomart image is:
+--[ RSA 2048 ]-----+
|      .      .      |
|    + + *   o      |
|  + + B = o .      |
|    o . + = . .    |
|    . . S + .      |
|    . . . . .      |
|    E   o          |
|    . o            |
|    . .            |
+--[ MD5 ]-----+
$ ls -l $HOME/.cge
total 8
-rw----- 1 username group 1679 Jan  6 11:49 id_rsa
-rw-r--r-- 1 username group 391 Jan  6 11:49 id_rsa.pub
```

**NOTE:** While this example shows creation of keys under `$HOME/.cge`, it can be used to place keys in any directory. If `$HOME/.cge` is not a convenient place to put the keys, follow the above procedure to generate the keys in some other (suitably protected) directory, then use the `--configDir` option to `cge-launch` or the `$CGE_CONFIG_DIR_NAME` environment variable to point to that directory. If it is required to use CGE specific keys that are stored on HDFS, create them in a temporary directory using this procedure, then copy them onto HDFS in the location of your choice (appropriately).

protecting them). Then use an HDFS URL as the value of `$CGE_CONFIG_DIR_NAME` or the argument to the `--configDir` option to the `cge-launch` command to select that directory instead of `$HOME/.cge` as the key directory.

Once this has been done, CGE will use the keys in the `.cge` directory instead of the ones in the `.ssh` directory and there should be no further problems with pass phrases.

## 6.3 Grant Other Users Access to Their CGE Query Engine

The Cray Graph Engine (CGE) can protect the contents of user-owned data sets from view/modification by unauthorized users via CGE instances that you run. Regardless of this protection, it is required to protect the raw data in user-owned data sets using traditional Linux file protection, otherwise users who have access to their data can start their query engine, using their data without knowledge. To ensure that only authorized users gain access to user-owned data, it is best to set the permissions on each directory containing a data set to permit access (read, write and execute/search) only by its owner, and then to set the permissions on the files in the directory to permit access (read and write) only to their owners.

As the owner of a running instance of a CGE, it is possible to control the list of users to whom access is granted. There are two modes of granting access to other users:

- Access to a single data set
- Access to any provided data set

A key first step to any of this is protecting owned data sets from being used under some other user's instance of CGE. If a user can run her own instance of CGE using your data, then you have no further control. So, if it is required to control access to owned data sets, make sure they are protected against access by users other than you. By setting the permissions on the data directory for the data set to `rwX-----` you achieve this by preventing other users from looking in that directory for files. If other users can be allowed to run their own instances of CGE using user-owned data, these permissions may be set any way desired.

Assuming data sets have been protected against other users, now individual users can be granted access. Regardless of whether you want to grant access to one or all data sets, you need the contents of each user's public key file from that user's `.ssh` directory. The user can follow the steps for setting up keys shown above if she does not have them yet. It is okay for the user to send you the public key(s) via e-mail, or any other method (including letting you copy them from the files yourself). They need to be appended to an appropriate `authorized_keys` file.

**IMPORTANT:** Remember that any user trying to connect with the server will need to authenticate the server as described in [Grant Basic Access to Owned Query Engines](#) or configure the CLI to trust Host Keys as described in [Eliminating the Interactive Host Key Verification](#).

Ask users to do the following after granting them access:

```
$ cge echo --db-port=73737
The authenticity of host localhost' can't be established.
RSA key fingerprint is d2:b4:ad:70:f1:44:d3:8a:f5:16:db:db:76:07:19:47.
Are you sure you want to continue connecting? [Yes/No]: yes
13835 [main] WARN com.cray.cge.communications.client.ssh.LoggingBridge - Permanently added 'localhost' (RSA) to the
list of known hosts.
14110 [main] INFO com.cray.cge.sparql.cli.lightweight.commands.debug.EchoCommand - Sending echo request...
14157 [main] INFO com.cray.cge.sparql.cli.lightweight.commands.debug.EchoCommand - Echoed data received and validated
successfully
```

**NOTE:** It is important to note that a user should NEVER add another user's public key to the user-owned `$HOME/.ssh/authorized_keys` file. Doing so will allow the user to login as the user who owns that file.

In the following example, it is assumed that `/lus/scratch/username/lubm0` directory contains one of user-owned data sets:

```
$ ls -ld /lus/scratch/username/lubm0
drwxr-xr-x 2 username group 4096 Oct 20 14:23 /lus/scratch/username/lubm0
$ chmod og-rwx /lus/scratch/username/lubm0
$ ls -ld /lus/scratch/username/lubm0
drwx----- 2 username group 4096 Oct 20 14:23 /lus/scratch/username/lubm0
$ ls -l /lus/scratch/username/lubm0/
total 4796
-rw-r--r-- 1 username group      221 Jan  6 13:13 authorized_keys
-rwxr-xr-x 1 username group 3321856 Oct  9 11:52 dbQuads
-rwxr-xr-x 1 username group 1568768 Oct  9 11:52 string_table_chars
-rw-r--r-- 1 username group    8192 Oct  9 11:52 string_table_chars.index
$ chmod og-rwx /lus/scratch/username/lubm0/*
$ ls -l /lus/scratch/username/lubm0/
total 4796
-rw----- 1 username group      221 Jan  6 13:13 authorized_keys
-rw----- 1 username group 3321856 Oct  9 11:52 dbQuads
-rw----- 1 username group 1568768 Oct  9 11:52 string_table_chars
-rw----- 1 username group    8192 Oct  9 11:52 string_table_chars.index
```

Now this data set can only be used by instances of the query engine that the user starts. Other users wanting access will need to connect with a client and will be subject to client authentication.

### 6.3.1 Grant Other Users Access to One of the Owned Data Sets

To grant a user access to one of your data sets, all you need to do is put the user's public key in the `authorized_keys` file in the same directory where your data set resides, as shown in the following example:

```
$ ls -l /lus/scratch/username/lubm0/
total 4792
-rwxr-xr-x 1 username group 3321856 Oct  9 11:52 dbQuads
-rwxr-xr-x 1 username group 1568768 Oct  9 11:52 string_table_chars
-rw-r--r-- 1 username group    8192 Oct  9 11:52 string_table_chars.index
$ cat my_friend_id_rsa.pub >> /lus/scratch/username/lubm0/authorized_keys
$ ls -l /lus/scratch/username/lubm0/
total 4796
-rw-r--r-- 1 username group      221 Jan  6 13:13 authorized_keys
-rwxr-xr-x 1 username group 3321856 Oct  9 11:52 dbQuads
-rwxr-xr-x 1 username group 1568768 Oct  9 11:52 string_table_chars
-rw-r--r-- 1 username group    8192 Oct  9 11:52 string_table_chars.index
$ cat /lus/scratch/username/built_lubm0/authorized_keys
ssh-rsa AAAAB3NzaC1yc2EAAAABIWAAAIEAxp7+CpYHL44jmuWeGXEMy+ijE/
X72f70YL8neITsR5gotXCiZh9V0G9ar8mNDlkoshN7Jp1qiRrQjYnY93hs9BBCz9kA5V9PhGC59qypEhNovYRo48lsUvTmHK0RWOVLfIZKNCKLVmbQubmEz
M0FfUoY/
ifNbTfrV4yGH2PNA4k= my_friend@myhost
```

Once you have done this, the user 'my\_friend' will have access to this data set only and not to all of your data sets. You can copy the `authorized_keys` file to any other data set you want to grant access to, and edit it as needed.

### 6.3.2 Grant Other Users Access to All of the Owned Data Sets

If you do not need to restrict access to specific data sets to a particular user, it is simpler to grant that user access to all of your data sets in one `authorized_keys` file. CGE uses a directory located at `$HOME/.cge` that allows you to set up configuration files that apply to all of your data sets. Users can grant access to all of their data sets by creating an `authorized_keys` file in this directory and putting authorized public keys in that file, as shown in the following example:

```
% mkdir -p $HOME/.cge
$ chmod o-w,g-w $HOME/.cge
$ cat my_friend_id_rsa.pub >> $HOME/.cge/authorized_keys
$ ls -l $HOME/.cge
total 4796
-rw-r--r-- 1 username group      221 Jan  6 13:13 authorized_keys
$ cat $HOME/.cge/authorized_keys
ssh-rsa AAAAB3NzaC1yc2EAAAABIWAAAIEAxp7+CpYHL44jmuWeGXEMy+ijE/
X72f70YL8neITsR5gotXCiZh9V0G9ar8mNDlkoshN7Jp1qiRrQjYnY93hs9BBCz9kA5V9PhGC59qypEhNovYRo48lsUvTmHK0RWOVLfIZKNCKLVmbQubmEz
M0FfUoY/
ifNbTfrV4yGH2PNA4k= my_friend@myhost
```

While this example shows placement of the global `authorized_keys` file in `$HOME/.cge`, it can be used to place the `authorized_keys` file in any directory. If `$HOME/.cge` is not a convenient place to put the

`authorized_keys` file, follow the above procedure to place it in some other (suitably protected) directory, then use the `--configDir` option to `cge-launch` or the `$CGE_CONFIG_DIR_NAME` environment variable to point to that directory. If a global `authorized_keys` file needs to be stored on HDFS, create the file using this procedure, then copy it onto HDFS in the location of choice (appropriately protecting it). Then use an HDFS URL as the value of `$CGE_CONFIG_DIR_NAME` or the argument to the `--configDir` option to `cge-launch` to select that directory instead of `$HOME/.cge`.

Now the user `my_friend` will have access to all of your data sets.

## 7 Built-in Graph Functions

---

SPARQL is intrinsically designed to find explicit patterns in graphs, using the basic graph patterns called out in SPARQL specifications. Often these patterns themselves create a graph that needs to be analyzed in a way that is not easily implemented with SPARQL's basic graph patterns. One example of this in the Lehigh University Benchmark (LUBM) ontology would be to find students who take courses from their advisers, and then find the shortest path through a social network between specific pairs of those students. Another example is to use betweenness centrality to find the most "central" (i.e., connecting the most entities not otherwise connected) entities in a graph, often a social network.

To address this other type of processing, CGE's SPARQL implementation has been extended to incorporate graph-function capability. This means that the input to the graph function is a graph, not just a few scalars, such as numbers or IRIs. This capability includes both the syntax that enables calling of graph functions, and a small number of built-in graph functions (BGFs) that are callable by any CGE user.

The built-in graph functions included in this release of CGE are:

- **BadRank:** Assigns a "badness" score to all vertices in the graph based on their nearness to known bad vertices.
- **Betweenness Centrality:** Ranks each vertex by how frequently it is on the shortest path between vertices.
- **Page Rank:** Measures the relative importance of a vertex in a graph.
- **Label Propagation:** Detects communities in networks and assigns vertices in the graph to communities.
- **S-T Connectivity:** Finds the shortest path, if one exists, between two vertices in the graph.
- **S-T Set Connectivity:** Finds the shortest path, if one exists, between a set of vertices designated as sources and a set of vertices designated as targets.
- **Triangle Counting:** Counts the total number of triangles in a graph.
- **Triangle Finding:** Finds all the triangles in the graph.
- **Vertex Triangle Counting:** Gathers statistics on the vertices based on the triangles they participate in and for non-cyclic triangles, their position in the triangle.

### 7.1 Combining Graph Algorithms with SPARQL

CGE provides an infrastructure for calling graph algorithms from within SPARQL queries. A graph algorithm is called via a CGE-specific SPARQL operator named `INVOKE`.

It is useful to note the following items:

1. The `INVOKE` operator cites the name of the graph algorithm being invoked, using an URI notation that is similar to that used for representing built-in functions in SPARQL.
2. Scalar arguments can be input to the graph algorithm via a parenthesized argument list.

3. The `INVOKE` clause is always preceded by a SPARQL `CONSTRUCT` clause, whose function in this context is to build the graph that is input to the graph algorithm. CGE provides the capability of nesting a `CONSTRUCT/INVOKE` clause within a `SELECT/WHERE` clause. This enables a subquery within a SPARQL query to select or produce a subgraph, which is used as input to the graph algorithm.
4. The `INVOKE` clause is immediately followed by a `PRODUCING` clause, whose function is to bind the results of the graph algorithm to specific SPARQL variables.
5. While RDF graphs may define many different types of subjects and objects, the CGE graph algorithms treat them all as homogeneous vertices and do not distinguish between them according to type, with the exception of functions that explicitly expect some vertices to be distinguished.
6. The `CONSTRUCT-INVOKE-PRODUCING` combination needs to be nested within a `SELECT-WHERE` clause.
7. For all CGE-specific built-in graph functions, if the query writer wants to specify a non-default value for an argument, values for the preceding arguments also need to be specified, even if default values for those arguments are to be used.

## 7.2 Invocation of a Graph Function

Four SPARQL constructs are involved while invoking graph functions. These include:

- `CONSTRUCT`
- `INVOKE`
- `PRODUCING`
- `SELECT-WHERE`

### 7.2.1 The `CONSTRUCT` Clause

There are three main differences between a standard SPARQL `CONSTRUCT` clause and the way it is used in CGE in a `CONSTRUCT-INVOKE-PRODUCING` combination. These differences are described below:

1. As mentioned above, the `CONSTRUCT-INVOKE-PRODUCING` combination always appears nested within the `WHERE` clause of a `SELECT` query.
2. While a standard SPARQL `CONSTRUCT` query returns an RDF graph to the user, the `CONSTRUCT` clause of a `CONSTRUCT-INVOKE-PRODUCING` combination does not return anything to the user; instead the constructed graph is passed to the graph algorithm as input, and then discarded after the graph algorithm completes execution.
3. Because the output of the nested `CONSTRUCT` clause is eventually discarded, CGE relaxes some of the rules for constructing RDF graphs. In particular, since some graph algorithms expect weighted edges, CGE allows predicates to be literals inside a nested `CONSTRUCT` clause.

### 7.2.2 The `INVOKE` Clause

In CGE, graph functions are invoked using the CGE-specific `INVOKE` keyword with the `CONSTRUCT` query form. The syntax of the `INVOKE` keyword is shown below:

```
INVOKE <http://cray.com/graphAlgorithm.graph_function> (arguments)
```



In the above example, *graph\_function* is the name of the graph function to be invoked and *arguments* is a comma-separated list of arguments to be provided to the graph function. The types and number of arguments in this list are dependent on the function being invoked.

## Using the **INVOKE** Keyword

```
SELECT *
  WHERE {
    CONSTRUCT {
      ?s ?p ?o .
    } WHERE {
      ?s ?p ?o .
    }
    INVOKE <http://cray.com/graphAlgorithm.graph_function> (42,0.19,"string")
    PRODUCING ?varX ?varY
  }
```

In the above example, the **INVOKE** keyword is used to invoke a graph function named "graph\_function" with three scalar arguments as well as the graph produced by the **CONSTRUCT** clause.

### 7.2.3 The **PRODUCING** Clause

In the Cray Graph Engine (CGE), the invocation of a graph function results in an intermediate result set. Ultimately, this is what enables graph functions to be composed with other SPARQL operators such as **UNION**, **ORDER BY**, or **FILTER**, as they also output an intermediate result set. The **PRODUCING** keyword can be used to bind the columns of the returned intermediate result set to SPARQL variables. The **PRODUCING** keyword accepts a list of SPARQL variable names which will be bound to the columns of the intermediate result set returned by the **INVOKE** keyword. Therefore, while using the **PRODUCING** keyword, it is required to know the following:

- How many columns will exist in the returned intermediate result set
- What set of values each column represents

The syntax of the **PRODUCING** keyword is shown below:

```
PRODUCING ?varA ?varB
```

In the above statement, *?varA* and *?varB* are the variables specified by the **PRODUCING** operator to be bound to columns of the returned vectors of results.

## Using the **PRODUCING** Clause

The community detection algorithm returns two columns of information. Information contained in these columns is described below:

- The first column contains each of the vertex IDs of the graph that was sent to the algorithm.
- The corresponding entry in the second column contains an integer that represents the identity of the community to which that vertex was assigned.

Thus the **PRODUCING** clause would specify variables that the query author chose to reflect the two vectors of data being returned, as shown in the following query snippet:

```
...
INVOKE <http://cray.com/graphAlgorithm.community>( )
```

```
PRODUCING ?vertexID ?communityID
...
```

## 7.3 Inputs to the Graph Function

Three types of inputs to a graph algorithm are possible:

1. The graph itself – Each graph function expects input to come from the output of the preceding `CONSTRUCT` operator.
2. Scalar inputs – Scalar values can be passed to the graph algorithm via a parenthesized list in the `INVOKE` clause.
3. Vector inputs – Sets of values can be input to the graph algorithm by adding them to the graph that the `CONSTRUCT` operator builds. Generally these inputs are distinguished in the input graph by a triple with a type predicate and a special type object.

In the following example, the Bad Rank algorithm expects to receive a set of vertex IDs of vertices considered to be spam, i.e, it could represent some other undesirable attribute. Note that the `WHERE` clause associated with the `CONSTRUCT` clause includes a `VALUES` clause, that names a set of vertices that are to be considered spam by the Bad Rank algorithm. That set of vertices is added to the `CONSTRUCT` clause's graph as a set of triples with a `rdf:type` predicate and the special object `cray:spamNode`. The scalar argument list of the `INVOKE` clause also specifies that this `cray:spamNode` object is to be used for identifying spam vertices. Similarly, a vector input to the graph algorithm can already be present in the database.

### Using Vector Inputs for Graph Algorithm

```
PREFIX cray: <http://cray.com/>
SELECT ?vertex ?ranking
{
  CONSTRUCT{
    ?sub ?pred ?obj .
    ?badNode a cray:spamNode .
  }
  WHERE {
    {
      ?sub ?pred ?obj .
    } UNION {
      VALUES ?badNode {
        <http://www.Department5.University0.edu/Course34>
        <http://www.Department6.University0.edu/GraduateCourse34>
        <http://www.Department14.University0.edu/GraduateCourse31>
        <http://www.Department5.University0.edu/Course34>
        <http://www.Department10.University0.edu/GraduateCourse25>
        <http://www.Department11.University0.edu/Course11>
        <http://www.Department13.University0.edu/GraduateStudent87>
      }
    }
  }
  INVOKE cray:graphAlgorithm.badrang (0.0001, .84, 0.01, cray:spamNode)
  PRODUCING ?vertex ?ranking
}
ORDER BY DESC (?ranking)
LIMIT 100
```

The above example shows the invocation of the Bad Rank algorithm with a set of spam vertices present in the input graph.

## 7.4 Sequence of Operators

The `PRODUCING` operator needs to immediately follow the `INVOKE` operator, which in turn needs to immediately follow the `WHERE` clause containing the `CONSTRUCT` operator. The `CONSTRUCT-INVOK-PRODUCING` combination should always appear as a nested subquery inside a `SELECT` clause's associated `WHERE` clause. Graph algorithms, like `SELECT` clauses themselves, can be nested arbitrarily deep in a query. Hence the sequence of operators that are involved in calling a graph algorithm is:

1. `CONSTRUCT-WHERE`
2. `INVOKE`
3. `PRODUCING`
4. `SELECT-WHERE`

**NOTE:** As mentioned earlier, the graph that is created by the `CONSTRUCT` clause that is part of a `CONSTRUCT-INVOK-PRODUCING` combination is never produced as output of the query; it is thrown away after it is used as input to the graph algorithm. If you want to see the graph that this `CONSTRUCT` clause builds, you must write a separate `CONSTRUCT` query.

### Example: Sequence of Operators

The following example illustrates the use of both spam and non-spam vertices with Bad Rank:

```
PREFIX cray: <http://cray.com/>
SELECT ?vertex ?ranking {
  CONSTRUCT {
    ?sub ?pred ?obj .
  } WHERE {
    {
      ?sub <http://bgf/isLinked> ?obj .
      ?sub <http://bgf/hasWeightLink> ?weightURI .
      ?obj <http://bgf/hasWeightLink> ?weightURI .
      ?weightURI <http://bgf/hasWeight> ?pred
    } UNION {
      ?sub <http://bgf/hasClassification> <http://bgf/spam> .
      BIND (<http://bgf/hasClassification> as ?pred) .
      BIND (<http://bgf/spam> as ?obj)
    } UNION {
      ?sub <http://bgf/hasClassification> <http://bgf/nospam> .
      BIND (<http://bgf/hasClassification> as ?pred) .
      BIND (<http://bgf/nospam> as ?obj)
    }
  }
  INVOKE cray:graphAlgorithm.badrnk (0.0001, .84, 0.01,
    <http://bgf/spam>, <http://bgf/nospam>, <http://bgf/hasClassification>)
  PRODUCING ?vertex ?ranking
}
ORDER BY DESC (?ranking)
LIMIT 100
```

## 7.5 Bad Rank

### URI

```
<http://cray.com/graphAlgorithm.badrank>
```

### Inputs and Default Values

- **Scalar inputs** - None.
- **Vector inputs** - The Bad Rank algorithm accepts input designation of a set of vertices that are known to be "bad", and optionally a set of vertices that are known to be "trusted", within the main input graph. This code can take three scalar arguments of type double.

Input	Default Value
The threshold of the maximum difference between per-vertex Bad Rank results from successive iterations of the algorithm below, which the algorithm will terminate.	0.0001
The probability that the next step in a (random) walk will be followed.	0.84
The probability that a random walk will take a next step to a bad vertex.	0.01
The URI that designates the object field of a triple that identifies a spam vertex	<http://cray.com/spamVertex>
The URI that designates the object field of a triple that identifies a non-spam, or trusted vertex.	<http://cray.com/nonspamVertex>
The URI that designates the predicate field of a triple that identifies either a spam or a non-spam vertex.	Defaults to the standard RDFS type predicate, <http://www.w3.org/1999/02/22-rdf-syntax-ns#type> The above can be abbreviated in a SPARQL query as "a".
The indicator that specifies whether or not normalization should be applied to results. Acceptable values for this parameter are 0 and 1.	1. If the default value is used, the scores are all mapped to floating point numbers between 0.0 and 1.0, with the maximum value found mapping to 1.0, the minimum score found mapping to 0.0, and other scores mapping between those values proportionately. If the value is set to 0, results will not be normalized and will be presented as Bad Rank computed them.

### Outputs

Bad Rank produces a two-column intermediate result that can be thought of as a set of pairs. The first item in each pair is the identifier of a vertex, whereas the second is the double-precision Bad Rank value of the vertex.

## 7.6 Betweenness Centrality

### URI and scalar arguments

```
<http://cray.com/graphAlgorithm.betweenness_centrality> (st_vx_ct, normalize)
```

In the above URI, *st\_vx\_ct* and *normalize* are used as examples.

### Description

This is the CGE-specific implementation of the classical vertex-betweenness-centrality algorithm. This algorithm assigns each vertex a numerical score. Take a given vertex *V*. In full generality, its betweenness score is defined to be the sum (over all other pairs of vertices) of the ratio of the number of shortest paths between that pair that go through *V*, over the total number of shortest paths between that pair. Thus it measures a sort of “importance” of each vertex, in terms of the shortest paths to other vertices that pass through it.

### Inputs and Default Values

- **Vector inputs** - None
- **Scalar inputs** - None
- **Input Graph** - The input graph to the Betweenness centrality function is expected to contain triples of the form (vertex1, weight, vertex2) where weight is an integer. The following table describes the inputs that may be provided to the `INVOKE` keyword to invoke the Betweenness Centrality function.

Parameter	Description	Default Value
<i>st_vx_ct</i>	<p>The <i>st_vx_ct</i> parameter can either be an integer or a decimal.</p> <ul style="list-style-type: none"> <li>• If the <i>starting_vertex_ctl</i> parameter is an integer, it represents how many starting vertices should be used when approximating the betweenness score of every vertex in the graph.</li> <li>• If the <i>starting_vertex_ctl</i> parameter is a decimal, it should be between 0.0 and 1.0. If a decimal argument is used, the decimal value will represent the fraction of the graph's vertices, randomly chosen, that will be used as starting vertices for approximating the betweenness scores. A value of 1.0 (the default) specifies that every vertex in the graph will be used as a starting vertex.</li> </ul>	1.0
<i>normalize</i>	The <i>normalize</i> parameter specifies whether or not the betweenness scores should be normalized. The acceptable values for this parameter are 0 and 1, where 1 specifies that betweenness scores should be normalized.	1

Parameter	Description	Default Value
	Normalizing the scores means to subtract from the betweenness score of each vertex the minimum betweenness score and then divide that partial result by the difference between the maximum and minimum betweenness scores found among all the vertices. Normalized scores will be between 0.0 and 1.0.	

## Outputs

A call to the Betweenness Centrality function returns a two-column intermediate result set. The first column contains the vertex identifier (URI), whereas the second column contains the centrality score of the vertex. In other words, each row of the output result set pairs a vertex's ID with a double-precision floating-point value representing the centrality score for that vertex.

## Example: Betweenness Centrality

```
PREFIX cray: <http://cray.com/>
SELECT ?vertices ?scores
WHERE {
  CONSTRUCT {
    ?sub ?pred ?obj .
  } WHERE {
    ?sub ?pred ?obj .
  }
  INVOKE cray:graphAlgorithm.betweenness_centrality(.01,1)
  PRODUCING ?vertices ?scores
}
ORDER BY DESC(?scores)
```

## 7.7 Community Detection Label Propagation (LP)

### URI

```
<http://cray.com/graphAlgorithm.community_detection_LP>
```

### Description

The Label Propagation algorithm is used for detecting communities in networks and assigns vertices in the graph to communities. Each vertex is initially assigned to its own community. At every step, each vertex looks at the community affiliation of all its neighbors, and updates their state to the mode community affiliation. The mode community affiliation takes into account the edge weights.

The Label Propagation algorithm is relatively inexpensive, but convergence is not guaranteed.

## Inputs and Default Values

The input graph to the Label Propagation function is expected to contain triples of the form (*vertex1*, *weight*, *vertex2*), where *weight* is an integer.

Input	Default Value
The number of steps that the algorithm executes. Currently an early exit is not included if convergence is detected. Therefore, the algorithm executes the number of steps specified in the input.	20

## Outputs

A call to the Label Propagation function returns an array of vertex IDs paired with an array of community IDs. These IDs can be used to identify which community each vertex was assigned to.

## Example: Label Propagation

```
PREFIX cray: <http://cray.com/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

SELECT ?vertex ?comm
WHERE{
  CONSTRUCT {
    ?sub ?weight ?obj .
  } WHERE {
    ?sub <http://wga/isLinked> ?obj .
    ?sub <http://wga/hasWeightLink> ?weightURI .
    ?obj <http://wga/hasWeightLink> ?weightURI .
    ?weightURI <http://wga/hasWeight> ?weight
  }
  INVOKE cray:graphAlgorithm.community_detection_LP(5)
  PRODUCING ?vertex ?comm
}
ORDER BY ?comm
```

## 7.8 Page Rank

### URI

```
<http://cray.com/graphAlgorithm.pagerank>
```

### Description

Page Rank can be used to measure the relative importance of a vertex in a graph.

### Inputs and Default Values

- **Vector inputs** - None.
- **Scalar inputs** - This algorithm accepts the following scalar inputs:

- `delta` - This is the convergence parameter. The default value of this parameter is 0.00001
- `dampen` - This is the dampening factor. The default value of this parameter is 0.8

OutputsPage Rank produces a two-column intermediate result that can be thought of as a set of pairs. The first item in each pair is the identifier of a vertex, whereas the second is the double-precision Page Rank value of the vertex.

### Example: Page Rank

The following example selects all of the edges from the default graph and calls S-T Set Connectivity on the resulting graph.

```
PREFIX cray: <http://cray.com/>

SELECT ?vertices ?pagerank
WHERE {
  CONSTRUCT{
    ?sub ?pred ?obj .
  }
  WHERE{
    { ?sub ?pred ?obj . }
  }
  INVOKE cray:graphAlgorithm.pagerank(0.0005,0.85)
  PRODUCING ?vertices ?pagerank
}
ORDER BY DESC(?pagerank)
```

## 7.9 S-T (Source – Target) Connectivity

### URI

```
<http://cray.com/graphAlgorithm.st_connectivity>
```

### Description

The S-T Connectivity function calculates the length of the path between two vertices, if one exists.

### Inputs and Default Values

- **Vector inputs** - None.
- **Scalar inputs** - The input graph to the S-T Connectivity function is expected to contain triples of the form (vertex1, predicate, vertex2) where the value of predicate is ignored. The S-T Connectivity function requires two scalar input arguments, which are the IRIs of the two vertices under consideration, source and target, respectively. This is illustrated in the example below:

```
INVOKE <http://cray.com/graphAlgorithm.st_connectivity>
(<urn:mySourceVertex>, <urn:myTargetVertex>)
```

In the above example, `<urn:mySourceVertex>` and `<urn:myTargetVertex>` are the IRIs of the source and target vertices, respectively.



## Outputs

The following example culls needed edges from the default graph and calls S-T Connectivity on the resulting graph.

### Example: S-T (Source Target) Connectivity

```
PREFIX cray: <http://cray.com/>

SELECT ?nHops
WHERE {
  CONSTRUCT {
    ?v1 ?p ?v2 .
  } WHERE {
    SELECT ?v1 ?v2 ?p
    WHERE {
      ?v1 <urn:hasLink> ?v2 .
      BIND(<urn:path> AS ?p)
    }
  }
  INVOKE cray:graphAlgorithm.st_connectivity(<http://ga.org/string#000/vertex#00000001>,
    <http://ga.org/string#000/vertex#00200000>)
  PRODUCING ?nHops
}
```

## 7.10 S-T Set Connectivity

### URI

```
<http://cray.com/graphAlgorithm.st_set_connectivity>
```

### Inputs and Default Values

- **Scalar inputs** - None.
- **Vector inputs** - The S-T Set Connectivity function accepts input of a set of vertices designated as sources and a set of vertices designated as targets. These are added to the constructed graph using the `<http://cray.com/sourceVertex>` and `<http://cray.com/targetVertex>` URIs, as well as the standard RDFS predicate `<http://www.w3.org/19990222-rdf-syntax-ns#type>`, which may be abbreviated as “a” in a SPARQL query.

Subject	Predicate	Object
Source vertex identifier	a	<code>&lt;http://cray.com/sourceVertex&gt;</code>
Target vertex identifier	a	<code>&lt;http://cray.com/targetVertex&gt;</code>

## Outputs

A call to the S-T Set Connectivity function returns an intermediate result set containing a single integer. The values and meaning of this integer are described below:

- If the integer's value is 0, there is no path between any pair of vertices with the source vertex taken from the source set and the target vertex taken from the target set.
- If the value is greater than 0, it represents the number of hops in the shortest path between any such pair of vertices.

**IMPORTANT:** The S-T Set Connectivity function will return an error in the following cases:

1. Nonexistence of input source and/or target vertex
2. Invalid input source and/or target vertex
3. Nonexistence of input source and/or target vertex in the input edge list

### Example: S-T Set Connectivity

The following example selects all of the edges from the default graph and calls S-T Set Connectivity on the resulting graph.

```
PREFIX cray: <http://cray.com/>
SELECT ?distance
WHERE {
  CONSTRUCT{
    ?sub ?pred ?obj .
    ?srcNode a cray:sourceVertex .
    ?trgNode a cray:targetVertex .
  }
  WHERE{
    {
      ?sub ?pred ?obj .
    }
    UNION {
      VALUES ?srcNode
      {
        <http://bgf.org/c/03/i/000000>
        <http://bgf.org/c/05/i/000000>
        <http://bgf.org/c/08/i/000003>
      }
    }
    UNION {
      VALUES ?trgNode
      {
        <http://bgf.org/c/05/i/000001>
        <http://bgf.org/c/08/i/000007>
      }
    }
  }
  INVOKE cray:graphAlgorithm.st_set_connectivity()
  PRODUCING ?distance
}
```

## 7.11 Triangle Counting

### URI

```
<http://cray.com/graphAlgorithm.triangle_counting>
```

### Description

Triangle Counting is used to count the total number of triangles in a graph.

### Inputs and Default Values

- **Vector inputs** - None.
- **Scalar inputs** - This algorithm accepts a single integer scalar argument. The value of this integer ranges from 0 to 4 and specifies which types of triangles are to be included in the count.
  - 0: Return a count of all the triangles in the graph, both cyclic (including rotations) and non-cyclic triangles
  - 1: Return a count of all the unique triangles in the graph, both cyclic and non-cyclic triangles
  - 2: Return a count of only the non-cyclic triangles
  - 3: Return a count of only the cyclic triangles (including rotations)
  - 4: Return a count of only the unique cyclic triangles

### Outputs

This algorithm returns a single integer containing the number of triangles.

### Example: Triangle Counting

```
PREFIX cray: <http://cray.com/>

SELECT ?total_num_triangles
WHERE {
  CONSTRUCT{
    ?sub ?pred ?obj .
  }
  WHERE{
    ?sub ?pred ?obj .
  }
  INVOKE cray:graphAlgorithm.triangle_counting(1)
  PRODUCING ?total_num_triangles
```

## 7.12 Vertex Triangle Counting

### URI

```
<http://cray.com/graphAlgorithm.vertex_triangle_counting>
```

## Description

The Vertex Triangle Counting algorithm is used to gather statistics on the vertices based on the triangles they participate in and for non-cyclic triangles, their position in the triangle.

## Inputs and Default Values

- **Vector inputs**- None.
- **Scalar inputs** - This algorithm accepts a single integer scalar argument. The value of this integer ranges from 0 to 4 and specifies which types of triangles are to be included in the counting statistics.
  - 0: Return a count of all the triangles in the graph, both cyclic (including rotations) and non-cyclic triangles.
  - 1: Return a count of all the unique triangles in the graph, both cyclic and non-cyclic triangles
  - 2: Return a count of only the non-cyclic triangles
  - 3: Return a count of only the cyclic triangles (including rotations)
  - 4: Return a count of only the unique cyclic triangles

## Outputs

Output is a four-column IRA. Each row in the IRA contains a vertex URI followed by a total count of the triangles for which it participates as either a `through_vertex`, `in_vertex`, or `out_vertex`, respectively. The `PRODUCING` clause should be interpreted as "`vertexID`", "`through_count`", "`in_count`", "`out_count`", where the counts refer to the number of triangles in which the vertex participates in that role.

## Example: Vertex Triangle Counting

```
PREFIX cray: <http://cray.com/>
SELECT ?id ?through ?in ?out
WHERE {
  CONSTRUCT{
    ?sub ?pred ?obj .
  }
  WHERE{
    ?sub ?pred ?obj .
  }
  INVOKE cray:graphAlgorithm.vertex_triangle_counting(0)
  PRODUCING ?id ?through ?in ?out
}
```

## 7.13 Triangle Finding

### URI

```
<http://cray.com/graphAlgorithm.triangle_finding>
```

## Description

The Triangle Finding algorithm is used to find all the triangles in the graph. The output can be customized to return either all triangles, or only the cyclic or non-cyclic triangles. The number of triangles in a given region of a graph is a good indicator of the density of that part of the graph.

## Inputs and Default Values

- **Vector inputs**- None.
- **Scalar inputs** - This algorithm accepts a single integer scalar argument. The value of this integer ranges from 0 to 4 and specifies which types of triangles are to be output..
  - 0 - Return all the triangles in the graph, both cyclic (including rotations) and non-cyclic triangles
  - 1: Return all the unique triangles in the graph, both cyclic and non-cyclic triangles
  - 2: Return only the non-cyclic triangles
  - 3: Return only the cyclic triangles (including rotations)
  - 4: Return only the unique cyclic triangles

## Outputs

The code returns a four-column IRA . Each row in the IRA represents the three URIs of the vertices of a triangle followed by a cyclic flag (set to 1 for cyclic, 0 for non-cyclic). The non-cyclic triangles are written out in the order of through\_vertex, in\_vertex, out\_vertex. The cyclic flag is considered optional in the `PRODUCING` clause in the case where only the URIs of the vertices are needed.

## Example: Triangle Finding

```
PREFIX cray: <http://cray.com/>

SELECT ?vertexID1 ?vertexID2 ?vertexID3 ?cyc
WHERE {
  CONSTRUCT{
    ?sub ?pred ?obj .
  }
  WHERE{
    ?sub ?pred ?obj .
  }
  INVOKE cray:graphAlgorithm.triangle_finding(1)
  PRODUCING ?vertexID1 ?vertexID2 ?vertexID3 ?cyc
```

## 8 Cray Graph Engine (CGE) Extension Functions

---

CGE provides a number of extension functions, including:

- Interval analytics functions. See [Cray Graph Engine \(CGE\) Interval Analytics Functions](#) on page 78 for more information.
- Haversine functions. See [Cray Graph Engine \(CGE\) Haversine Functions](#) on page 82 for more information.
- Square root function. See [Cray Graph Engine \(CGE\) Square Root Function](#) on page 82 for more information.

### 8.1 Cray Graph Engine (CGE) Interval Analytics Functions

#### Intervals

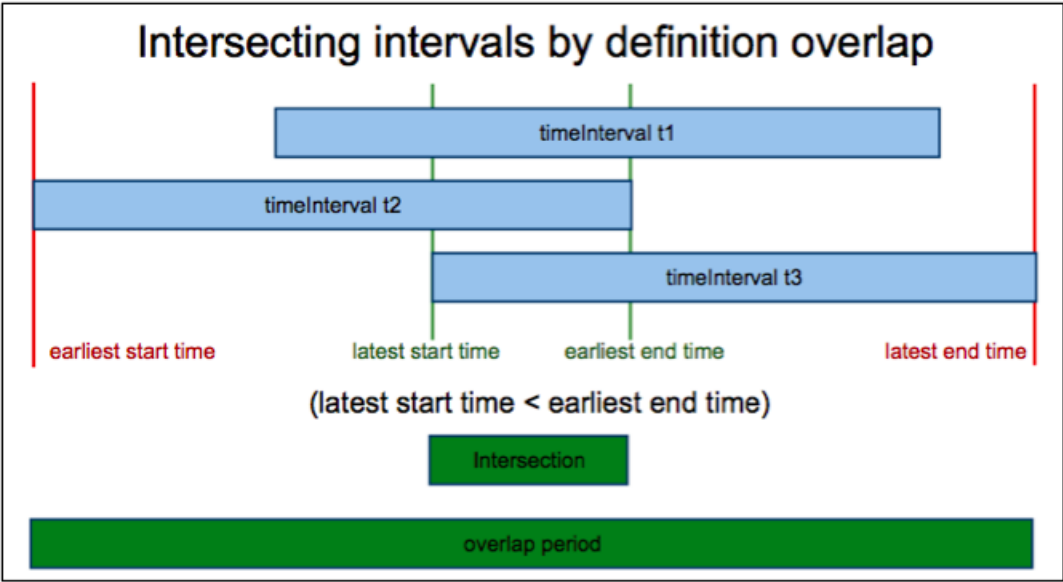
An interval is defined as the sequence between any two variables of compatible atomic types, where one defines the start of the interval and the other defines the end of the interval. The interval is inclusive of the start and end.

#### Intersecting and Non-Intersecting Intervals

Two or more intervals are said to be intersecting if there is an interval in time during which all of the intervals under consideration are present. More precisely, intersecting time intervals are those where the latest start time is less than the earliest end time, i.e., no period has ended before the last period to begin, has started. This period is termed as an intersection and starts at the beginning of the last interval to start and ends at the end of the first interval to end.

On the other hand, non-intersecting time intervals are those where a period has ended before the last period to begin has started, as illustrated in the following figure:

Figure 9. Intersecting Intervals



### Continuous and Non-Continuous Intervals

Two or more intervals are said to be “continuous” if there is at least one interval present during the complete span, from the start of the first interval to start, to end of the last interval to end. In non-continuous intervals, at least one gap (period within which no intervals are present) is present between the intervals under consideration. This is illustrated in the following figure:

Figure 10. Continuous Period With Non-Intersecting Intervals

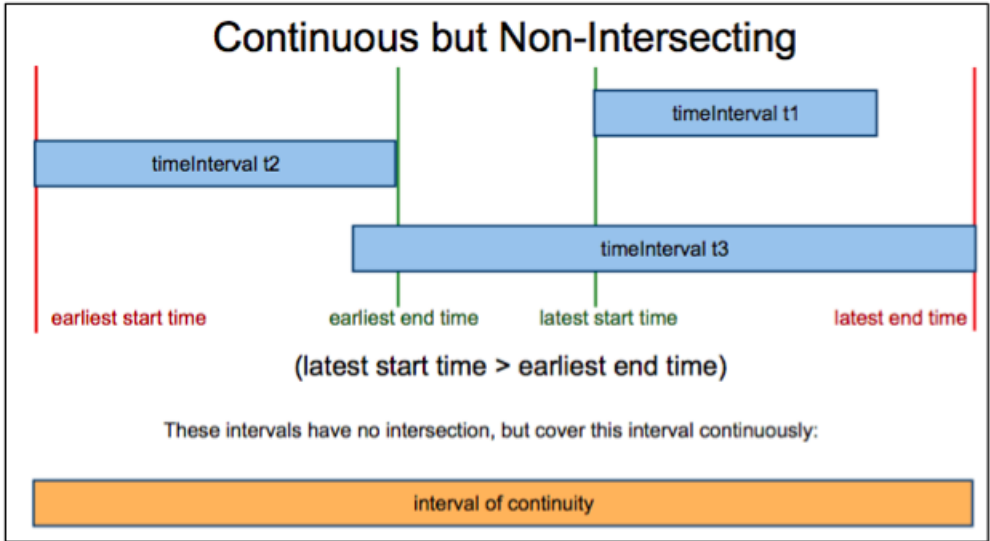
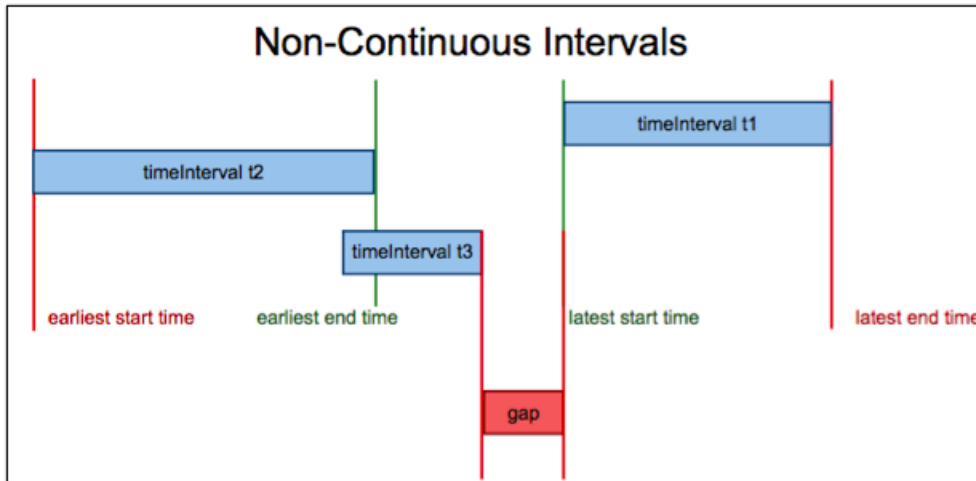


Figure 11. Non-Continuous Period with Non-Intersecting Intervals



## Using Interval Analytics Functions For Temporal Analysis

Interval functions can be used to gather fine-grained detail about intervals. For example, they can be used to:

- Determine if a time period that ends at the same time is contiguous with one that starts at the same time.
- Determine whether or not two or more time intervals intersect.
- Determine the continuity of a given time period.

## Function Prefix

The prefix to use when using interval functions in queries is:

PREFIX arq: <<http://jena.hpl.hp.com/ARQ/function#>>

## List of Interval Analytics Functions

**NOTE:** The names of all the CGE interval functions are case-sensitive. Interval functions work with any type that has a < comparison, e.g., numerics and strings.

The list of CGE-specific interval analytics functions, their syntax and description is provided in the following table:

Table 12. List of CGE-specific Interval Functions

Function	Description
<code>listmin(element1, .... elementN)</code>	This function returns the smallest item in the comma-separated list of items provided as arguments.
<code>listmax(element1, .... elementN)</code>	This function returns the largest item in the list of arguments.
<code>iscontinuous(start1,end1,... startN, endN)</code>	This is a pairwise function that accepts a list of comma-separated list of start and end times and determines whether or not there is a gap between the intervals under consideration.



Function	Description
	<ul style="list-style-type: none"> <li>• True when there is complete coverage from earliest starting time to latest end time, i.e. there are no gaps in the coverage.</li> <li>• False if there is any gap in the coverage</li> </ul>
<code>isintersecting(start1, end1, .... startN, endN)</code>	<p>This is a pairwise function that determines whether or not there is a period within which all the intervals under consideration are present. This function returns:</p> <ul style="list-style-type: none"> <li>• True when there is an interval where all intervals are present.</li> <li>• False if there is no interval when all intervals are present</li> </ul>
<code>duration(startTime, endTime)</code>	<p>This function uses the Unix epoch and time functions to calculate the duration between the start and end times, which are provided as arguments. This function returns the <code>xsd:dayTimeDuration</code> between <code>startTime</code> and <code>endTime</code>.</p> <p><b>NOTE:</b> This function only accepts dates starting from July 5, 1776.</p>

**NOTE:** Although the `listmin()`, `listmax()`, `iscontinuous()` and `isintersecting()` functions support all SPARQL compatible types, the arguments provided to these function should all be of compatible atomic types, otherwise an `xsd_error` will be returned. Furthermore, the `duration()` function will return an `xsd_error` in the following cases:

- Either of the arguments are not of type `xsd:dateTime`
- The sum of (`duration( xsdDate1, xsdDateTime2)` - `duration( xsdDateTime2, xsdDate1)`) will not be zero. This is because `xsdDate` is defined to span 24 hours (for standard days), and it is assumed that the start time is at the beginning of the day, and the end time is at the end of the day

There are a few important items to note when using the interval analytics functions:

- The interval analytic functions do not fully support the `xsd:date` and `xsd:time` data types and may return incorrect results; users should avoid these two types.
- Comparisons of `xsd:date` and `xsd:dateTime` within the same day may return unexpected results. `xsd:date` and `xsd:dateTime` comparisons are supported outside of the 14 hour time zone range and the 24 hour day span of `xsd:date`.
- `xsd:date` results are now included when filtering on `xsd:dateTime` (outside the same day) and vice versa (`xsd:dateTime` results when filter on `xsd:date`). If strict `xsd:dateTime` results (or `xsd:date` results) are required, the appropriate data type filter should be added.

- The `duration()` function supports combinations of `xsd:date` and `xsd:dateTime`. If an `xsd:date` result is the start time, the duration will start at the beginning of the day. Similarly, if the `xsd:date` result is the end time, the duration will end at the end of the day.

## 8.2 Cray Graph Engine (CGE) Haversine Functions

CGE supports the `haversinometers()` and `haversinemiles()` functions to enable support for spatially aware applications. These functions are based on the Haversine formula, which is an equation that calculates the great-circle distance between two points on a sphere from the longitudes and latitudes of the two points. For more information, visit [http://en.wikipedia.org/wiki/Haversine\\_formula](http://en.wikipedia.org/wiki/Haversine_formula).

The syntax of CGE Haversine functions is shown below:

- `afq:haversinometers(latStart, longStart, latEnd, longEnd)`
- `afq:haversinemiles(latStart, longStart, latEnd, longEnd)`

**NOTE:** The `haversinometers()` and `haversinemiles()` functions are case sensitive.

### Inputs

Both the CGE `haversinometers()` and `haversinemiles()` functions accept the following inputs in `xsd:decimal`, `xsd:double` and `xsd:float` formats:

- **atStart** – The starting position of the latitude (dimensions of the values in degrees)
- **longStart** – The starting position of the longitude (dimensions of the values in degrees)
- **latEnd** – The ending position of the latitude (dimensions of the values in degrees)
- **longEnd** – The ending position of the longitude (dimensions of the values in degrees)

Acceptable latitude values range from  $-90$  to  $90$ , whereas acceptable longitude values range from  $-180$  degrees to  $180$  degrees.

**NOTE:** Important: The functions will return an empty value if:

- Invalid position coordinates are provided
- Empty input values are provided
- Insufficient parameters are provided.

### Output

The `haversinometers()` function returns the distance between two points in meters, whereas the `haversinemiles()` function returns the distance between two points in miles.

### Function Prefix

The prefix to use when using CGE Haversine functions in queries is:

PREFIX `afq: <http://jena.hpl.hp.com/ARQ/function#>`

## 8.3 Cray Graph Engine (CGE) Square Root Function

The square root function, `sqrt()` is used to retrieve the square root of the specified number

### Syntax

The syntax of the square root function is:

```
sqrt(argument)
```

**NOTE:** The name of the `sqrt()` function is case sensitive.

### Function Prefix

The prefix to use when using the `sqrt()` function in queries is:

```
PREFIX afn: <http://jena.hpl.hp.com/ARQ/function#>
```

Table 13. CGE Square Root Function's Examples

Argument Type	Example
Integer	<pre>PREFIX afn: &lt;http://jena.hpl.hp.com/ARQ/function#&gt; SELECT ?a { BIND (afn:sqrt("9223372036854775807"^^ &lt;http://www.w3.org/2001/XMLSchema#integer&gt;) AS ?a) }</pre>
Decimal	<pre>PREFIX afn: &lt;http://jena.hpl.hp.com/ARQ/function#&gt; SELECT ?a { BIND (afn:sqrt(4294967296.0) AS ?a) }</pre>
Float	<pre>PREFIX afn: &lt;http://jena.hpl.hp.com/ARQ/function#&gt; SELECT ?a { BIND (afn:sqrt ("3.4E38"^^xsd:float) AS ?a) }</pre>
Double	<pre>PREFIX afn: &lt;http://jena.hpl.hp.com/ARQ/function#&gt; SELECT ?a { BIND (afn:sqrt("1.797E308"^^xsd:double) AS ?a) }</pre>
Boolean	<pre>PREFIX afn: &lt;http://jena.hpl.hp.com/ARQ/function#&gt; SELECT ?a { BIND (afn:sqrt(true) AS ?a) }</pre> <p><b>NOTE:</b> Passing “true” as the Boolean argument returns 1, whereas passing “false” as the Boolean argument returns 0.</p>

**NOTE:** The `sqrt()` function will return an empty value if a negative number is provided as an argument. Furthermore, the `sqrt()` function will return an empty value if arguments of certain types are used. These argument types include:

- `xsd:dateTime`
- String
- IRI
- Arbitrary data type

You can also use derived data types as arguments to the `sqrt()` function, as shown in the following query:

```
PREFIX afn: <http://jena.hpl.hp.com/ARQ/function#>
SELECT ?a{ BIND (afn:sqrt ("18446744073709551615"^^<http://www.w3.org/2001/XMLSchema#positiveInteger>) AS ?a) }
```

**NOTE:** Executing the `sqrt()` function when a negative derived type is used as an argument will result in an empty value.

## 9 Cray Graph Engine (CGE) Property Path Support

CGE does not natively support all the SPARQL 1.1 property paths features, however it does support certain types of property paths.

**NOTE:** CGE's property path support should be used with care. This support is disabled by default and must be explicitly enabled by the user. Contact Cray Support for additional information.

- **Simple Property Paths** - By default, simple property paths that are equivalent to simple fixed length Basic Graph Patterns (BGPs) are supported. This means that property paths consisting of only the sequence / and inverse ^ operators are permitted, since these can be written out as a simple BGP using blank node variables. For example:

```
SELECT * WHERE
{
  ?s <urn:a>/<urn:b> ?o
}
```

Can be rewritten as follows:

```
SELECT * WHERE
{
  ?s <urn:a> _:p0 .
  _:p0 <urn:b> ?o .
}
```

- **Complex Property Paths Emulation** - Some more complex property paths can be emulated through query rewriting, which expands the property paths into an equivalent query form.

**NOTE:** It is important to be aware that this support is only emulation, and may not provide complete answers that a SPARQL engine with native property path support would produce.

The following table details the additional operators, which may be emulated and the restrictions and limitations on that emulation.

*Table 14. Additional Operators that May be Emulated*

Operator	Example	Description	Additional Notes
*	?s <urn:a>* ?o	Finds paths of zero or more steps between two nodes in the graph	<ul style="list-style-type: none"> <li>• Path to which the * operator applied <b>must</b> be either a predicate or inverse predicate</li> <li>• Evaluating the zero length portion of the path may be very expensive particularly if both variables are unbound</li> <li>• Paths are evaluated only up to a maximum length (default 5) which</li> </ul>

Operator	Example	Description	Additional Notes
			<p>may be user configured on a per-query basis</p> <ul style="list-style-type: none"> <li>Expands into a UNION that looks for paths of each length up to the specified maximum</li> </ul>
+	<code>?s &lt;urn:a&gt;+ ?o</code>	Finds paths of one or more steps between two nodes in the graph	<ul style="list-style-type: none"> <li>Path to which the + operator applied <b>must</b> be either a predicate or inverse predicate</li> <li>Paths are evaluated only up to a maximum length (default being 5) which may be user configured on a per-query basis</li> <li>Expands into a UNION that looks for paths of each length up to the specified maximum</li> </ul>
?	<code>?s &lt;urn:a&gt;? ?o</code>	Finds paths of zero or one steps between two nodes in the graph	<ul style="list-style-type: none"> <li>Path to which the ? operator applied <b>must</b> be either a predicate or inverse predicate</li> <li>Evaluating the zero length portion of the path may be very expensive particularly if both variables are unbound</li> <li>Expands into a UNION that looks for paths of length zero and one</li> </ul>
	<code>?s &lt;urn:a&gt;   &lt;urn:b&gt; ?o</code>	Finds paths between two nodes that use any of the alternative paths given	<ul style="list-style-type: none"> <li>Paths to which the   operator applied may themselves be complex but only paths that are predicates or inverse predicates are guaranteed to expand into a valid query</li> <li>Expands into a UNION that considers each alternative, where the alternative is itself a property path it may be further expanded as necessary</li> </ul>
! ( property )	<code>?s ! &lt;urn:a&gt; ?o</code>	Find paths between two nodes that <b>do not</b> pass through a given predicate	<ul style="list-style-type: none"> <li>The negated property set operator only applies to predicates or inverse predicates and thus can always be expanded</li> <li>Expands into a MINUS that considers all paths and then eliminates the undesirable paths</li> </ul>

## Enabling Emulation

CGE also provides the option to change the maximum length of paths (for the expansion of the \* and + operators), as shown in the following example:

```
% cge-cli query --opt-on optPathExpand --path-expansion 3 paths.rq
```

The above query would run the query with property path expansion enabled and a maximum path length of 3.

**NOTE:** This value can be set to any desired value, however it is important to note that the higher this value is set to, the more complex the query that will be generated. This will result in slower performance because the database server will need to search for longer paths. Therefore, it is recommended to set the length of paths to the minimum possible value for optimal emulation performance. It is also important to note that setting a maximum length of zero or less will result in disabling the expansion.

## 10 Cray Graph Engine (CGE) Quick Reference

The order in which CGE operations should be performed is:

### Step 1: Set up SSH keys

If the following command allows re-logging into the login node without a password, then the SSH keys are set up sufficiently for using CGE.

```
$ ssh localhost
```

On the other hand, if the previous command fails and there are existing SSH keys that do not use pass-phrases or have the `ssh-agent` defined, then try the following:

```
$ cat ~/.ssh/id_*.pub >> ~/.ssh/authorized_keys
```

At this point, if it is possible to run the aforementioned test and to re-log in to the login node without using a password, pass-phrase, or `ssh-agent`, then this step can be considered to be complete. If, on the other hand, the aforementioned test fails, there are no SSH keys defined yet, the following commands can be used to set them up:



**CAUTION:** Ensure that there are no existing SSH keys because this will overwrite any existing keys. Also, do not specify a pass-phrase when running `ssh-keygen`

```
$ mkdir -p ~/.ssh
$ chmod 700 ~/.ssh
$ ssh-keygen
$ chmod 600 ~/.ssh/id_*
$ chmod 600 ~/.ssh/authorized_keys
```

If the existing SSH key(s) use pass-phrase(s) or the `ssh-agent`, or if a more complex SSH key configuration is required, see [Cray Graph Engine \(CGE\) Security Mechanisms](#) on page 56. This section also contains information about fine-tuning access to CGE instances.

### Step 2: Start the CGE Server

The `cge-launch` command launches the CGE query engine and enables creating and building a database in a single step.

The following is an example of using the `cge-launch` command:

```
$ cge-launch -o pathtoResultsDir -d path -l logfile
```

In the preceding example:

- `-o` - Specifies the path to a directory where you want the result files produced by queries to be placed.



**CAUTION:** This path **MUST** be a directory.

- `-d` - Specifies the path to the directory containing the data set to be loaded into the server. This directory must contain all input data files for the data set.

**NOTE:** This directory **MUST** contain at least one of the following if the data set is being built for the first time with CGE (only one of these will actually be used):

- `dataset.nt` - This file contains triples and must be named `dataset.nt`
  - `dataset.nq` - This file contains quads and must be named `dataset.nq`
  - `graph.info` - This file contains a list of pathnames or URLs to files containing triples or quads and must be named `graph.info`.
- `-l` - Specifies a log file to capture the command output from the run. If the database server is logging to `stderr`, this log file will capture that information as well. There are two special argument values for this: `':1'` and `':2'`, which refer to `stdout` and `stderr`, respectively, so that the log can be directed to either of those. If the `-l` option is specified, the `cge-launch` command runs silently, producing no output to the terminal `stdout/stderr`.

For more information, see [Launch the CGE Server Using the `cge-launch` Command](#) on page 20 and [Building a Database](#) on page 12.

### Step 3: Execute CGE CLI Commands (Optional)

CGE CLI commands can be executed after the CGE query engine has been launched. Following is an example of using the CGE `nvp-info` CLI command:

```
$ cge-cli nvp-info
```

CGE CLI features a number of commands, which are documented in the [Cray Graph Engine \(CGE\) Command Line Interface](#) on page 22 section.

### Step 3: Start up the CGE Front End Server to Connect with the CGE Server (Optional)

The CGE graphical user interface and SPARQL endpoints can be accessed once the database has been launched. This can be accomplished by launching the web server that provides the user interface on a login node of the system where CGE is running.

```
$ cge-cli fe --ping
```

The `--ping` option in the preceding example is used to verify that the database can be connected to immediately upon launch and that any failure is seen immediately. Not doing so may delay and hide failures. If the ping operation does not succeed, and it is certain that the user executing this command is the only user running CGE, and that everything else is set up correctly, the user should go back to the first step and make sure that the SSH keys are set up right. The system may prompt to trust the host key when the `fe` command is run for the first time. The default URL to access the UI is `http://<hostname>:3756/dataset`, where `hostname` is used as an example for the web server's name. For more information, see [Launch the CGE Web Server Using the `fe` Command](#) on page 34.

Alternatively, the following command can be used to have the web server continue running in the background with its logs redirected, even if disconnected from the terminal session:

```
$ nohup cge-cli fe > web-server.log 2>&1 &
```



## Step 4: Access and Use the CGE Front End (Optional)

For more information, see [Access the Cray Graph Engine \(CGE\) Graphical User Interface](#) on page 43.

## Shutdown the CGE Server

- Shut down the CGE server using the `shutdown` command, as shown in the following example:

```
$ cge-cli shutdown
```

For more information, see [Shutdown the CGE Server Using the shutdown Command](#) on page 41.

- Shut down the CGE front end if it was started.

## Additional Information

**Cancelling a query** - To cancel a query, hit CTRL-C on the window where the CGE server was launched or locate the CGE server instance's PID on the login node and use `kill -INT <PID>`. After that, re-launch CGE.

## 11 Use the Cray Graph Engine (CGE) for a Hello World Example

---

### About this task

In this example, a tiny RDF triples database and query are created in such a way that it creates a sort of "Hello World" output.

### Procedure

1. Create a `.nt` file. This is the original, readable representation of the database

```
<http://cray.com/example/spaceObject> <http://cray.com/example/hasName> "World" .
<http://cray.com/example/spaceObject> <http://cray.com/example/hasName> "Home Planet" .
<http://cray.com/example/spaceObject> <http://cray.com/example/hasName> "Earth" .
<http://cray.com/example/greeting> <http://cray.com/example/text> "Hello" .
<http://cray.com/example/greeting> <http://cray.com/example/text> "Hi" .
```

2. Store the `.nt` file in the directory that has been selected or created for it.

**IMPORTANT:** This directory must contain a file named `dataset.nt` if it contains triples or `dataset.nq` if it contains quads.

3. Select or create another directory into which the query engine should write the results and then launch the CGE server in a terminal window.

```
$ cge-launch -I 1 -N 1 -d /dirContainingExample/example -o \
/dirContainingExampleOutput -l :2
```

For more information about the `cge-launch` command, and its parameters, see the `cge-launch` man page or [Launch the CGE Server Using the cge-launch Command](#) on page 20.

The server will output a few pages of log messages, as it starts up and converts the database to its internal representation. When it finishes, the system will display a message similar to the following:

```
$ Serving queries on nid00057 16702
```

4. In another terminal window, launch the CGE front end:

```
$ cge-cli fe
```

When the CGE front end has been launched, a message similar to the following will be returned:

```
249 [main] INFO com.cray.cge.cli.commands.sparql.ServerCommand - CGE SPARQL Protocol Server has started and is
ready to accept HTTP requests on localhost:3756
```

**NOTE:** The CGE SPARQL protocol server listens at port 3756, which is the default port ID.

Now the browser can be started.

- Point the browser at the system the query engine is running on via (in the default case) port 3756. For example, if the query engine is running on a system named *my-machine*, point the browser at: `http://my-machine:3756/dataset/`. This will connect you to the CGE front end, and windows for editing queries or updates.

Figure 12. CGE Front End

Query Interface | Update Interface | Checkpoint Interface | Server Information | Edit Server Configuration

### Query Interface

SPARQL Query:

☐ Force text/plain as the response Content-Type (forces results to be displayed in browser)

Server NVPs:

# Enter NVPs one per line in properties file format e.g.  
 #cge.server.DoMemoryLeakDetection=1  
 #  
 # Lines beginning with a # are comments  
 # You can also enter ARQ optimizer flags in the following form e.g.

Server Logging Options:

Server Log Level: Use Server Default ⓘ

Server Log String:  (Printed on each server log line for this request)

☐ Disable all server logging for this request

- Execute a query against the dataset created by typing in the query and selecting the **Run Query** button.

Figure 13. CGE Hello World Query Example

Query Interface | Update Interface | Checkpoint Interface | Server Information | Edit Server Configuration

### Query Interface

SPARQL Query:

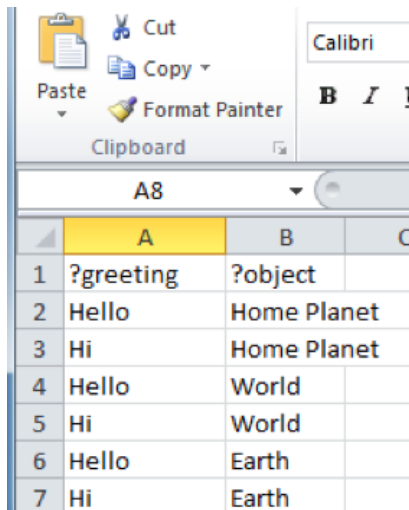
```
SELECT ?greeting ?object
WHERE {
  <http://cray.com/example/spaceObject> <http://cray.com/example/hasName> ?object .
  <http://cray.com/example/greeting> <http://cray.com/example/text> ?greeting .
}
```

After the query finishes executing, the output file containing the query's results will be stored in the output directory that was specified in the `cge-launch` command. This can be verified by listing the contents of that directory and reviewing the contents of the output file, as shown in the following example:

```
$ cd /dirContainingExampleOutput
$ ls
queryResults.34818.2015-10-05T19.33.53Z000.tsv
$ cat queryResults.34818.2015-10-05T19.33.53Z000.tsv
?greeting    ?object
"Hello"      "Home Planet"
"Hi"         "Home Planet"
"Hello"      "World"
"Hi"         "World"
"Hello"      "Earth"
"Hi"         "Earth"
```

Since the Excel application can read `.tsv` files, results can also be viewed in Excel, as shown in the following figure:

*Figure 14. Viewing CGE Output in Excel*



	A	B	C
1	?greeting	?object	
2	Hello	Home Planet	
3	Hi	Home Planet	
4	Hello	World	
5	Hi	World	
6	Hello	Earth	
7	Hi	Earth	

7. To halt the CGE server, execute:

```
$ cge-cli shutdown
```

## 12 CGE Java API

The CGE Java API consists of four Java JAR files and a ReadMe file:

1. `cge-java-api-vx.x.x-with-dependencies.jar` - Contains executable Java classes. Using the executable Java classes, users can write their own Java code to perform CGE actions like launching the server, querying, updating, checkpointing, etc.
2. `cge-java-api-examples-vx.x.x-sources.jar` - Contains sample source code. Users can build their own Java programs referring to these examples.
3. `cge-java-api-vx.x.x-javadocs.jar` - Contains documentation of Java classes that the CGE API is comprised of.
4. `cge-java-examples-vx.x.x-with-dependencies.jar` - Contains executable sample source code. Samples provided with the CGE API can also be executed, as their entry points are included in the executable sample source code
5. `ReadMe_JavaAPI.txt` - Contains commands and scripts documented in this guide. These commands and scripts are suitable for copying and pasting for execution.

Where `vx.x.x` is the version code, the initial version being `v1.0.0`.

This section describes how to utilize these `.jar` files.

### Extracting the source code

To extract the Java source code, run the following command:

```
$ jar -xf PATH_TO_JAR/cge-java-api-examples-v1.0.0-sources.jar
```

A directory structure similar to the following should appear:

```
./META-INF
./META-INF/MANIFEST.MF
./com
./com/cray
./com/cray/cge
./com/cray/cge/api
./com/cray/cge/api/examples
./com/cray/cge/api/examples/ExampleUtils.java
./com/cray/cge/api/examples/Checkpoint.java
./com/cray/cge/api/examples/LaunchAndShutdown.java
./com/cray/cge/api/examples/Shutdown.java
./com/cray/cge/api/examples/LaunchAndTerminateOnJvmExit.java
./com/cray/cge/api/examples/Status.java
./com/cray/cge/api/examples/ComprehensiveExample.java
./com/cray/cge/api/examples/UpdateExisting.java
./com/cray/cge/api/examples/IsRunning.java
./com/cray/cge/api/examples/StatusExisting.java
./com/cray/cge/api/examples/QueryExisting.java
```

```
./com/cray/cge/api/examples/CheckpointExisting.java
./com/cray/cge/api/examples/Update.java
./com/cray/cge/api/examples/Query.java
./com/cray/cge/api/examples/Config.java
./com/cray/cge/api/examples/hooks
./com/cray/cge/api/examples/hooks/InheritIOHook.java
./com/cray/cge/api/examples/LaunchOnly.java
```

The `.java` files are source-code examples that users can refer to when building their own Java programs. These are also included in the executable Java classes, so users can execute them directly if desired.

## Extracting the API Class Documentation

To extract the API class documentation, run the following command:

```
$ $ jar -xf PATH_TO_JAR/cge-java-api-v1.0.0-javadocs.jar
```

A directory structure similar to the following should appear. Please note that there is an extensive directory structure that exists under the `./com` directory, but it is not displayed in the following.

```
./META-INF
./META-INF/MANIFEST.MF
./resources
./resources/titlebar.gif
./resources/tab.gif
./resources/titlebar_end.gif
./resources/background.gif
./com
./allclasses-frame.html
./overview-frame.html
./overview-summary.html
./package-list
./deprecated-list.html
./serialized-form.html
./index.html
./help-doc.html
./index-all.html
./allclasses-noframe.html
./constant-values.html
./stylesheet.css
./overview-tree.html
```

This file system is meant to be run with a web-browser.

- To view on Windows systems, right-click on **overview-summary.html** and then select the **open with** menu option to open the file using the web-browser of interest, such as Internet Explorer or Chrome.
- To execute on Linux systems, ensure that an X-Windows server is running on the target computer, and that the `$DISPLAY` environment variable is set appropriately. Then execute with the web-browser of interest, for example using `firefox help-doc.html&`.

In either case, the documentation should appear for point-and-click viewing.

## Using API Executables

There are several ways to utilize the executable Java classes `.jar` files, as illustrated by the following:

- Use Java API via Maven

- Use Java API via Java Development Kit (JDK)
- Use Java API via pre-built `main` entry points
- Use Java API Comprehensive Example program

## 12.1 Use CGE Java API via Maven

### About this task

The following procedure illustrates a use-case for using Maven to develop a Java application program to execute the `cge-launch` command. This sample program will utilize the `LaunchOnly.java` sample file.

### Procedure

1. Create an application framework.

```
$ mvn archetype:generate -DgroupId=com.mycompany.launchonly \
-DartifactId=my-launchonly -DarchetypeArtifactId=maven-archetype-quickstart \
-DinteractiveMode=false
```

This is a standard Maven command to make a framework for developing a Java application; in this case, a framework is created in the new `my-launchonly` area.

2. Switch to the `my-launchonly` directory.

```
$ cd my-launchonly
```

3. Install executable classes as a local `jar` file

```
$ mvn install:install-file \
-Dfile=./cge-java-examples-v1.0.0-with-dependencies.jar \
-DgroupId=com.cray.cge.api \
-DartifactId=cge-user-apis \
-Dversion=1.0.0 \
-Dpackaging=jar
```

This is a standard Maven command for creating a directory structure for an executable `.jar` file.

4. Copy the `LaunchOnly.java` file into this area of the new `mycompany` directory.

```
$ cp ../LaunchOnly.java \
./src/main/java/com/mycompany/launchonly/LaunchOnly.java
```

5. Develop a `pom.xml` file for the application or overwrite the default `pom.xml` file. There can be many variations of this file, the following is shown as a suggestion for the contents of this file:

```
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/
2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/
maven-v4_0_0.xsd">
  <modelVersion>4.0.0</modelVersion>
  <groupId>com.cray.cge.api.examples</groupId>
```

```

<artifactId>my-launchonly</artifactId>
<packaging>jar</packaging>
<version>1.0.0</version>
<name>my-launchonly</name>
<url>http://maven.apache.org</url>
<build>
  <plugins>
    <plugin>
      <groupId>org.apache.maven.plugins</groupId>
      <artifactId>maven-assembly-plugin</artifactId>
      <version>2.6</version>
      <executions>
        <execution>
          <goals>
            <goal>attached</goal>
          </goals>
          <phase>package</phase>
          <configuration>
            <descriptorRefs>
              <descriptorRef>jar-with-dependencies</descriptorRef>
            </descriptorRefs>
            <archive>
              <manifest>
                <mainClass>com.cray.cge.api.examples.LaunchOnly</mainClass>
              </manifest>
            </archive>
          </configuration>
        </execution>
      </executions>
    </plugin>
  </plugins>
</build>
<dependencies>
  <dependency>
    <groupId>junit</groupId>
    <artifactId>junit</artifactId>
    <version>3.8.1</version>
    <scope>test</scope>
  </dependency>
  <dependency>
    <groupId>com.cray.cge.api</groupId>
    <artifactId>cge-user-apis</artifactId>
    <version>1.0.0</version>
  </dependency>
</dependencies>
</project>

```

## 6. Build the package.

Users can first delete all files in their `~/ .m2` directory, as Maven will download what it needs to build the package into this area.

```
$ mvn clean package
```

## 7. Execute the code using one of the following:

- `$ java -jar ./target/my-launchonly-1.0.0-jar-with-dependencies.jar /path-to-dataset ./`
- `$ java -cp ./target/my-launchonly-1.0.0-jar-with-dependencies.jar com.cray.cge.api.examples.LaunchOnly /path-to-dataset ./`



- `$ export CLASSPATH=./target/my-launchonly-1.0.0-jar-with-dependencies.jar;`  
`java com.cray.cge.api.examples.LaunchOnly /path-to-dataset ./`

Note the following items:

- The `./` argument specifies where the program will place the result files.
- `/path-to-dataset` is the directory path to the user's dataset area,  
e.g. `/lus/scratch/ripple/mkdb/sp2b/25k`
- Outputs are the `cge_launcher.log` and `cge_runtime.log` files.

Outputs to `stdout` from execution should be similar to the following:

```
Launcher arguments are:
cge-launch -d /path-to-dataset -o my-launchonly/. -l cge_runtime.log -p
12345 --nodeCount 1 --imagesPerNode 1 --sessionTimeout 900
Starting CGE...
CGE not yet ready (1 seconds elapsed)
CGE not yet ready (2 seconds elapsed)
CGE not yet ready (3 seconds elapsed)
CGE ready in 7 seconds
CGE is running
```

## 8. Shut down the CGE CLI

```
$ cge-cli shutdown --db-port 22334
```

## 12.2 Use CGE API via Java Development Kit (JDK)

### About this task

This procedure illustrates a use-case where programmers utilize the JDK directly for Java program development.

### Procedure

#### 1. Create a program

```
$ mkdir jdk_is_running
$ cd jdk_is_running
```

This sample program will determine if the CGE server is running.

#### 2. Create a directory to put the `.jar` files into and move to that directory

```
$ mkdir cray_jars
```

#### 3. Copy the `cge-java-examples-v1.0.0-with-dependencies.jar` file to this area.

```
$ cp /opt/cray/cge/default/lib/java/cge-java-examples-v1.0.0-with-
dependencies.jar cray_jars
```

#### 4. Create a file called `Manifest.txt`

```
$ touch Manifest.txt
```

5. Edit the `Manifest.txt` file to contain the following lines.

```
Main-Class: com.cray.cge.api.examples.IsRunning
Class-Path: cray_jars/cge-java-examples-v1.0.0-with-dependencies.jar
<blank_line>
```

**IMPORTANT:** The file containing the above lines must have a blank line (press Enter/Return to create a blank line) the end of the file. in the preceding example, `<blank_line>` is used to indicate a line with no characters.

6. Make a directory to locate the source file

```
$ mkdir -p com/cray/cge/api/examples
$ cp IsRunning.java com/cray/cge/api/examples
```

7. Copy the `IsRunning.java` sample code into the new directory.

**NOTE:** The directory name must match the package name in the source code.

See 'Extracting the Source Code' for the location of `IsRunning.java`.

8. Build the package

```
$ javac -classpath cray_jars/cge-java-examples-v1.0.0-with-dependencies.jar \
com/cray/cge/api/examples/IsRunning.java
```

9. Run via the Java interpreter

```
$ java -cp cray_jars/cge-java-examples-v1.0.0-with-dependencies.jar: \ com/
cray/cge/api/examples/IsRunning
```

The output from execution will indicate that either CGE is running or CGE is not running .

10. Build executable JAR file

```
$ jar cvfm0 out.jar Manifest.txt com/cray/cge/api/examples/IsRunning.class
```

11. Run the executable JAR file

```
$ java -jar out.jar
```

Output from execution should be either `CGE is running` or `CGE is not running`.

## 12.3 Use CGE API via Pre-built Main Entry Points

### About this task

This procedure illustrates a use-case where developers can run the main entry points in the `cge-java-examples-v1.0.0-with-dependencies.jar` file directly. These correspond with the example Java files containing the source line `public static void main(String[] args)`, of which there are

several, for example: `./com/cray/cge/api/examples/Shutdown.java`. See [Extracting the source code](#) on page 93 for the location of `Shutdown.java`.

From any directory that contains `cge-java-examples-v1.0.0-with-dependencies.jar`. The following examples show actual paths to datasets and repositories.

## Procedure

### 1. Launch the CGE server

```
$ java -cp cge-java-examples-v1.0.0-with-dependencies.jar \
com.cray.cge.api.examples.LaunchOnly /lus/scratch/ripple/mkdb/sp2b/25k ./
```

**NOTE:** The `./` argument specifies where the program will place the result files. In this example, the `/lus/scratch...` area contains a typical `sp2b` test dataset.

Output includes the files `cge_launcher.log` and `cge_runtime.log`. Outputs that appear on stdout should be similar to the following:

```
Launcher arguments are:
cge-launch -d /lus/scratch/ripple/mkdb/sp2b/25k \
-o /ufs/home/users/$USER/my_repository/. \
-l /ufs/home/users/$USER/my_repository/cge_runtime.log \
-p 22334 --nodeCount 1 --imagesPerNode 1 --sessionTimeout 900
Starting CGE...
CGE not yet ready (1 seconds elapsed)
CGE not yet ready (2 seconds elapsed)
CGE ready in 5 seconds
CGE is running
```

### 2. Check if the CGE server is running

```
$ java -cp cge-java-examples-v1.0.0-with-dependencies.jar
com.cray.cge.api.examples.IsRunning
```

The following will be displayed on stdout:

```
CGE is running
```

### 3. Execute a query

```
$ $ java -cp cge-java-examples-v1.0.0-with-dependencies.jar
com.cray.cge.api.examples.QueryExisting \
/home/users/$USER/cge-benchmark/cge_queries/sp2b/2.txt > query_results.out >
SELEC(COUNT(?s) as ?CNT) {?s ?p ?o}
```

In the preceding example, `2.txt` is a pre-defined query meant for the `sp2b-25k` dataset. It is also possible to create a `2.txt` file with `SELECT(COUNT(?s) as ?CNT) {?s ?p ?o}` as the only line.

The query results will by default go to stdout and a `.tsv` file. In this example, the query results are extensive so we redirect the default output to a file named `query_results.out`. The generated file `queryResults.2016-05-12T16.07.46Z000.12512.tsv` is also shown:

```
$ ls -l query_results.out
-rw-r--r-- 1root 292653May 1211:07query_results.out
$ ls -l queryResults.2016-05-12T16.07.46Z000.12512.tsv
-rw-r--r-- 1root 299154May 1211:07queryResults.2016-05-12T16.07.46Z000.12512.tsv
```

#### 4. Update

```
$ java -cp cge-java-examples-v1.0.0-with-dependencies.jar \
com.cray.cge.api.examples.UpdateExisting > updates.log
```

This update to the dataset is the simple default `INSERT DATA {<urn:s> <urn:p> <urn:o>}`, found in the `Update.java` sample file. See [Extracting the source code](#) on page 93 for the location of `Update.java`. An argument such as `.examples.UpdateExisting ./path-to-file` can be used to specify a more complex update command contained in a file.

The update output will by default go to stdout and consists of CGE log entries. In this example, the results are redirected to a file `updates.log`:

```
$ ls -l updates.log
-rw-r--r-- 1 root 8056 May 12 11:09 updates.log
```

#### 5. Create a checkpoint

```
$ java -cp cge-java-examples-v1.0.0-with-dependencies.jar \
com.cray.cge.api.examples.CheckpointExisting chkpt.sp2b.25k
```

This will checkpoint the dataset to a subdirectory in the dataset area, which in this example is named `./checkpoints/chkpt.sp2b.25k_Thu_May_12_12:00:10_CDT_2016`. Files in this directory consist of:

- `./checkpoints/chkpt.sp2b.25k_Thu_May_12_12:00:10_CDT_2016/string_table_chars.index`
- `./checkpoints/chkpt.sp2b.25k_Thu_May_12_12:00:10_CDT_2016/export_dataset.nq`
- `./checkpoints/chkpt.sp2b.25k_Thu_May_12_12:00:10_CDT_2016/string_table_chars`
- `./checkpoints/chkpt.sp2b.25k_Thu_May_12_12:00:10_CDT_2016/dbQuads`

The following will be displayed on stdout:

```
Checkpoint successful - see directory ./checkpoints
```

#### 6. Shutdown the CGE server

```
$ java -cp cge-java-examples-v1.0.0-with-dependencies.jar \
com.cray.cge.api.examples.Shutdown > shutdown.log
```

The shutdown output will by default go to stdout and consists of CGE log entries. In this example, the results are redirected to a file `shutdown.log`.

```
$ ls -l shutdown.log
-rw-r--r-- 1 root 663 May 12 12:05 shutdown.log
```

## 12.4 Use Case: A Comprehensive Java Program

### About this task

This procedure illustrates a use-case where Java programmers create a Java program that will execute several features of the Java API, namely:

1. Launching the CGE server

2. Running query and update commands
3. Checkpointing the dataset
4. Shutting down the CGE server.

This case utilizes the `cge-java-api-v1.0.0-with-dependencies.jar` file.

## Procedure

1. Create an application framework:

```
mvn archetype:generate \
-DgroupId=com.myapp.runcge \
-DartifactId=my-run-cge \
-DarchetypeArtifactId=maven-archetype-quickstart \
-DinteractiveMode=false
```

This is a standard Maven command to make a framework for developing a java application; in this case, a framework is created in the new `my-run-cge` directory area.

2. Install executable classes as a local JAR file

```
mvn install:install-file \
-Dfile=/opt/cray/cge/default/lib/java/cge-java-api-v1.0.0-with-dependencies.jar \
-DgroupId=com.cray.cge.api \
-DartifactId=cge-user-apis \
-Dversion=1.0.0 \
-Dpackaging=jar
```

This is a standard Maven command to install an executable `.jar` file. A directory structure similar to the following should appear in the user's `.m2/repository` directory:

```
./com
./com/cray
./com/cray/cge
./com/cray/cge/api
./com/cray/cge/api/cge-user-apis
./com/cray/cge/api/cge-user-apis/maven-metadata-local.xml
./com/cray/cge/api/cge-user-apis/1.0.0
./com/cray/cge/api/cge-user-apis/1.0.0/cge-user-apis-1.0.0.pom
./com/cray/cge/api/cge-user-apis/1.0.0/cge-user-apis-1.0.0.jar
```

3. Switch to `my-run-cge` directory and then copy the file to `./src/main/java/com/myapp/runcge`.

```
$ cd my-run-cge
$ cp ../com/cray/cge/api/examples/ComprehensiveExample.java \
./src/main/java/com/myapp/runcge
```

4. Write the Java source code. Following is the source code for the proposed Java program used in this example. This can be copied into the `./src/main/java/com/myapp/runcge` framework area:

```
package com.myapp.runcge;

// for standard java processing
import java.io.File;
import java.io.IOException;
```

```

import java.util.Collections;
import java.util.concurrent.TimeUnit;
import java.util.Date;
import org.apache.commons.lang3.StringUtils;

// for prepare cge launcher
import com.cray.cge.api.CgeConnection;
import com.cray.cge.api.CgeLauncher;
import com.cray.cge.api.builders.CgeConnectionBuilder;
import com.cray.cge.api.builders.CgeLauncherBuilder;
import com.cray.cge.api.builders.JobOptionsBuilder;
import com.cray.cge.communications.messaging.exceptions.CommunicationsException;

// for query execution
import org.apache.jena.atlas.io.IO;
import com.hp.hpl.jena.query.QueryExecution;
import com.hp.hpl.jena.query.ResultSet;
import com.hp.hpl.jena.query.ResultSetFormatter;
import com.hp.hpl.jena.sparql.resultset.ResultsFormat;
import com.cray.cge.sparql.engine.results.ResultsMetadata;

// for update execution
import com.hp.hpl.jena.update.UpdateProcessor;

// for log4j initialization
import org.apache.log4j.Level;

/**
 * Example that demonstrates launching CGE, run query, update, checkpoint and
 * shutdown.
 */
public class ComprehensiveExample
{
    /**
     * Default sparql commands run by this example
     */
    public static final String DEFAULT_QUERY = "SELECT ?type (COUNT(?s) AS ?
usages) WHERE { ?s a ?type } GROUP BY ?type";
    public static final String DEFAULT_UPDATE = "INSERT DATA { <urn:s> <urn:p>
<urn:o> }";

    // default runtime values the user can override with command line args.
    public static String dataset_area = "./";
    public static String output_area = "./";
    public static String checkpoint_area = "./";
    public static int NODE_COUNT=1;
    public static int IMAGE_COUNT=1;
    public static String query = DEFAULT_QUERY;
    public static String update = DEFAULT_UPDATE;
    public static String query_filename = null;
    public static String update_filename = null;
    public static File queryFile=null;
    public static File updateFile=null;
    public static int RUNTIME_TIMEOUT = 10; // minutes
    public static int STARTUP_TIMEOUT = 15; // seconds
    public static int CGE_CONNECTION_TIMEOUT = 3; // seconds
    public static int SERVER_PORT = 56789; // valid port number range: 1-65535
    (1-1023 may require superuser privileges)

    private static void showUsage()
    {

```

```

        System.out.println("\nExercises CGE by launching the server, run query,
update, checkpoint, and shutdown");
        System.out.println("Usage:");
        System.out.println("    -c          CGE server-connect timeout (seconds)
(default: 3)");
        System.out.println("    -d          Directory containing dataset
(default: ./)");
        System.out.println("    -k          Checkpoint dataset directory
(default: ./)");
        System.out.println("    -n          Number of nodes to run CGE on (default:
1)");
        System.out.println("    -i          Number of images to run CGE on (default:
1)");
        System.out.println("    -o          Outputs directory (created if does not
exist) (default: ./)");
        System.out.println("    -p          CGE server port (default: " + SERVER_PORT
+ ") range: 1024-65535 (1-1023 as su)");
        System.out.println("    -q          File with sparql query (default: '" +
DEFAULT_QUERY + "')");
        System.out.println("    -r          Runtime timeout (minutes) (default: 10)");
        System.out.println("    -s          Startup timeout (seconds) (default: 15)");
        System.out.println("    -u          File with sparql update (default: '" +
DEFAULT_UPDATE + "')");
        System.out.println("\n");
    }

    /**
     * Expects a next argument, prints an error and exists if none present
     * @param i Current Argument Index
     * @param argv Arguments
     * @param arg Current Argument for which we expect a value as the next
argument
     */
    private static void expectNextArg(int i, String[] argv, String arg)
    {
        if (i >= argv.length - 1)
        {
            System.err.println("Unexpected end of arguments, expected a value to
be specified after the " + arg + " option");
            System.exit(1);
        }
    }

    /**
     * Parses Arguments
     * @param argv Arguments
     */
    private static void parseArgs(String[] argv)
    {
        for (int i = 0; i < argv.length; i++)
        {
            try
            {
                String arg = argv[i];
                if (arg.equals("-h"))
                {
                    // Show Usage Summary and exit
                    showUsage();
                    System.exit(1);
                }
            }
            // we have one or more name-value-pair ("NVP") args

```

```

        // (e.g., "-n 4").
        expectNextArg(i, argv, arg); // exits if "value" of the NVP is
absent
        if (arg.equals("-d")) {
            dataset_area = argv[++i];
        } else if (arg.equals("-o")) {
            output_area = argv[++i];
        } else if (arg.equals("-k")) {
            checkpoint_area = argv[++i];
        } else if (arg.equals("-n")) {
            NODE_COUNT = Integer.parseInt(argv[++i]);
        } else if (arg.equals("-p")) {
            SERVER_PORT = Integer.parseInt(argv[++i]);
            if ((SERVER_PORT > 65535) || (SERVER_PORT < 1)) {
                System.out.println("Error: arg '-p server port " +
SERVER_PORT + "' value out of range");
                System.exit(1);
            } else if (SERVER_PORT < 1024) {
                System.out.println("Notice: arg '-p server port " +
SERVER_PORT + "' may require su privileges");
            }
        } else if (arg.equals("-c")) {
            CGE_CONNECTION_TIMEOUT = Integer.parseInt(argv[++i]);
        } else if (arg.equals("-r")) {
            RUNTIME_TIMEOUT = Integer.parseInt(argv[++i]);
        } else if (arg.equals("-i")) {
            IMAGE_COUNT = Integer.parseInt(argv[++i]);
        } else if (arg.equals("-q")) {
            query_filename = argv[++i];
            queryFile = new File(query_filename);
            if (!queryFile.isFile()) {
                System.out.println("Problem with " + query_filename + " -
does not exist or not a file");
                System.exit(1);
            }
        } else if (arg.equals("-u")) {
            update_filename = argv[++i];
            updateFile = new File(update_filename);
            if (!updateFile.isFile()) {
                System.out.println("Problem with " + update_filename + " -
does not exist or not a file");
                System.exit(1);
            }
        } else {
            System.err.println("Illegal Option " + arg);
            showUsage();
            System.exit(1);
        }
    }
}
catch (NumberFormatException numEx)
{
    //Occurs when a numeric parameter is expected but not received
    System.err.println("Illegal value '" + argv[i] + "' encountered
after option " + argv[i-1] + " when an integer value was expected");
    System.exit(1);
}
}
}

// main entry point

```



```

    public static void main(String[] args) throws IOException,
        CommunicationsException, InterruptedException {

        // suppress "log4j WARN" messages
        org.apache.log4j.Logger.getRootLogger().setLevel(org.apache.log4j.Level.OFF);

        parseArgs(args);

        String DB_LOG = "cge_runtime.log";
        String LAUNCHER_LOG = "cge_launcher.log";

        // show the runtime selections
        System.out.println("CgeLauncherBuilder - start ... ");
        System.out.println("... dataset          " + dataset_area);
        System.out.println("... output area       " + output_area);
        System.out.println("... checkpoint area   " + checkpoint_area);
        System.out.println("... query file        " + query_filename);
        System.out.println("... update file       " + update_filename);
        System.out.println("... node count        " + NODE_COUNT);
        System.out.println("... image count       " + IMAGE_COUNT);
        System.out.println("... server port       " + SERVER_PORT);
        System.out.println("... run timeout       " + RUNTIME_TIMEOUT);
        System.out.println("... start timeout     " + STARTUP_TIMEOUT);
        System.out.println("... connect timeout   " + CGE_CONNECTION_TIMEOUT);

        // Prepare the launcher
        CgeLauncher launcher = new CgeLauncherBuilder()
            .forExistingDatabase(dataset_area)
            .usingOutputDirectory(output_area)
            .usingDatabaseLogFile(DB_LOG)
            .usingLauncherLogFile(LAUNCHER_LOG)
            .onPort(SERVER_PORT)
            .withJobOptions(new JobOptionsBuilder()
                .withNodes(NODE_COUNT)
                .withImagesPerNode(IMAGE_COUNT)
                .withMaximumRuntime(RUNTIME_TIMEOUT)
                .build())
            .build();

        System.out.println("CgeConnectionBuilder - start ...");

        CgeConnection cge = new CgeConnectionBuilder()
            .usingLauncher(launcher)
            .withConnectionTimeout(CGES_CONNECTION_TIMEOUT,
TimeUnit.SECONDS)
            .onHost("localhost")
            .onPort(SERVER_PORT)
            .build();

        System.out.println("CgeConnectionBuilder - done!");

        // Start CGE
        startCge(cge, STARTUP_TIMEOUT, false);
        if (cge.isRunning()) {
            System.out.println("CGE is running");
        } else {
            System.err.println("CGE failed to start");
        }
    }

```

```

// run query
if (cge.isRunning()) {
    System.out.println("start query... ");
    if (queryFile != null) {
        query = IO.readWholeFileAsUTF8(queryFile.getAbsolutePath());
    }
    System.out.println("running query:\n\n" + query + "\n");

    ResultsMetadata results = cge.querySummary(query);
    if (results.isSuccessful()) {
        System.out.println("query complete - see results in " +
results.getLocation());
    } else {
        System.out.println("Error: query failed with: " +
results.getError());
    }
} else {
    System.err.println("CGE appears to not be running");
}

// run update
if (cge.isRunning()) {
    System.out.println("start update... ");
    if (updateFile != null) {
        update = IO.readWholeFileAsUTF8(updateFile.getAbsolutePath());
    }
    System.out.println("running update:\n\n" + update + "\n");
    // Updates are evaluated via the Apache Jena ARQ UpdateProcessor API
    UpdateProcessor up = cge.update(update);
    up.execute();
    System.out.println("update complete - see " + DB_LOG + " for log
entries.");
} else {
    System.err.println("CGE does not appear to be running");
}

// run checkpoint
if (cge.isRunning()) {
    System.out.println("start checkpoint... ");
    // Checkpoint the database currently in use by cge-server.
    Date curr_date = new Date();
    String clean_date = curr_date.toString();
    File cpDir = new File(checkpoint_area, "checkpoint" + File.separator +
dataset_area.replace('/', '_') + "_" + clean_date.replace(' ', '_'));
    cge.checkpoint(cpDir, true);
    System.out.println("Checkpoint successful - see directory " +
checkpoint_area + "/checkpoint");
} else {
    System.err.println("CGE failed to start");
}

// Shutdown
if (cge.isRunning()) {
    System.out.println("start shutdown... ");
    cge.stop();
    cge.getProcess().waitFor();
    System.out.println("...shutdown complete");
}

```

```

    }

    System.out.println( "exiting...");
    System.exit(0);
}

/**
 * Starts the CGE instance represented by the given connection
 *
 * @param cge
 *         CGE connection
 * @param maxWaitSeconds
 *         Maximum number of seconds to wait for start up
 * @param returnOnInterrupt
 *         Whether to return if interrupted while waiting
 * @throws IOException
 *         Thrown if there is a problem starting CGE
 */
public static void startCge(CgeConnection cge, int maxWaitSeconds, boolean
returnOnInterrupt) throws IOException {

    System.out.println("Starting CGE...");
    cge.start();

    long startTime = System.currentTimeMillis();
    while (TimeUnit.MILLISECONDS.toSeconds(System.currentTimeMillis() -
startTime) < maxWaitSeconds)
    {
        try
        {
            Thread.sleep(1000);
        }
        catch (InterruptedException e)
        {
            // Ignore or return as appropriate
            if (returnOnInterrupt)
            {
                System.out.println(String.format("Interrupted while waiting for
CGE to ready (%d seconds elapsed)",
TimeUnit.MILLISECONDS.toSeconds(System.currentTimeMillis() - startTime)));
                return;
            }
        }
        if (cge.isRunning())
        {
            System.out.println(String.format("CGE ready in %d seconds",
TimeUnit.MILLISECONDS.toSeconds(System.currentTimeMillis() - startTime)));
            return;
        }
        System.out.println(String.format("CGE not yet ready (%d seconds
elapsed)",
TimeUnit.MILLISECONDS.toSeconds(System.currentTimeMillis() - startTime)));
    }
}

```

5. Use the following `pom.xml` file, which is developed for building application code:

```
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/
2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/
maven-v4_0_0.xsd">
  <modelVersion>4.0.0</modelVersion>
  <groupId>com.myapp.runcge</groupId>
  <artifactId>my-run-cge</artifactId>
  <packaging>jar</packaging>
  <version>1.0.0</version>
  <name>my-run-cge</name>
  <url>http://maven.apache.org</url>
  <build>
    <plugins>
      <plugin>
        <groupId>org.apache.maven.plugins</groupId>
        <artifactId>maven-assembly-plugin</artifactId>
        <version>2.6</version>
        <executions>
          <execution>
            <goals>
              <goal>attached</goal>
            </goals>
            <phase>package</phase>
            <configuration>
              <descriptorRefs>
                <descriptorRef>jar-with-dependencies</descriptorRef>
              </descriptorRefs>
              <archive>
                <manifest>
                  <mainClass>com.myapp.runcge.ComprehensiveExample</mainClass>
                </manifest>
              </archive>
            </configuration>
          </execution>
        </executions>
      </plugin>
    </plugins>
  </build>
  <dependencies>
    <dependency>
      <groupId>junit</groupId>
      <artifactId>junit</artifactId>
      <version>3.8.1</version>
      <scope>test</scope>
    </dependency>
    <dependency>
      <groupId>com.cray.cge.api</groupId>
      <artifactId>cge-user-apis</artifactId>
      <version>1.0.0</version>
    </dependency>
  </dependencies>
</project>
```

6. Build

```
$ mvn clean package
```

Users can first delete all files in their `~/ .m2` directory, as Maven will download what it needs to build the package into this area.

7. Execute the `help` command. This will show the help menu for the application, defining the optional runtime arguments (note that all have default values):

```
system:~/comprehensive_test/my-run-cge> java -jar target/my-run-cge-1.0.0-jar-
with-dependencies.jar -h
Exercises CGE by launching the server, run query, update, checkpoint, and
shutdown
Usage:
  -c          CGE server-connect timeout (seconds) (default: 3)
  -d          Directory containing dataset (default: ./)
  -k          Checkpoint dataset directory (default: ./)
  -n          Number of nodes to run CGE on (default: 1)
  -i          Number of images to run CGE on (default: 1)
  -o          Outputs directory (created if does not exist) (default: ./)
  -p          CGE server port (default: 56789) range: 1024-65535 (1-1023 as su)
  -q          File with sparql query (default: 'SELECT ?type (COUNT(?s) AS ?
usages) WHERE { ?s a ?type } GROUP BY ?type')
  -r          Runtime timeout (minutes) (default: 10)
  -s          Startup timeout (seconds) (default: 15)
  -u          File with sparql update (default: 'INSERT DATA { <urn:s> <urn:p>
<urn:o> }')
```

8. Execute. The following is a sample execution command and resulting output:

```
system:~/comprehensive_test/my-run-cge> java -jar target/my-run-cge-1.0.0-jar-
with-dependencies.jar -d \
/lus/scratch/ripple/mkdb/sp2b/25k -p 12345 -r 2 -i 4 -n 3 -k /lus/scratch/
temp -q sp2b_query_9.txt
CgeLauncherBuilder - start ...
... dataset          /lus/scratch/ripple/mkdb/sp2b/25k
... output area      ./
... checkpoint area  /lus/scratch/temp
... query file       sp2b_query_9.txt
... update file      null
... node count       3
... image count      4
... server port      12345
... run timeout      2
... start timeout    15
... connect timeout  3
CgeConnectionBuilder - start ...
CgeConnectionBuilder - done!
Starting CGE...
CGE not yet ready (1 seconds elapsed)
CGE not yet ready (2 seconds elapsed)
CGE not yet ready (3 seconds elapsed)
CGE ready in 7 seconds
CGE is running
start query...
running query:

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
SELECT DISTINCT ?predicate
WHERE {
  {
```

```

    ?person rdf:type foaf:Person .
    ?subject ?predicate ?person
  } UNION {
    ?person rdf:type foaf:Person .
    ?person ?predicate ?object
  }
}

query complete - see results in file: /lus/scratch/comprehensive_test/my-run-
cge/queryResults.<YEAR-DAY-TIME>.tsv
start update...
running update:

INSERT DATA { <urn:s> <urn:p> <urn:o> }

update complete - see cge_runtime.log for log entries.
start checkpoint...
Checkpoint successful - see directory /lus/scratch/temp/checkpoint
start shutdown...
...shutdown complete
exiting...

```

The dataset being referenced in this example is `sp2b-25k`, and resides in the `/lus/scratch/ripple/mkdb/sp2b/25k` directory. Three nodes and four images were specified for the CGE server to execute on. Two minutes were allowed for the execution, 15 seconds for the startup, and three seconds to connect to the server, which should have been started at port 12345. The CGE server began executing within seven seconds of the initial command. The query being run was in the local file `sp2b_query_9.txt`, the content of which is shown following the "running query" banner. The output area for the query results was the local directory `./`. An update command file is not specified and so the default was used. The checkpointed dataset went into the `/lus/scratch/temp/checkpoint` directory under a sub-directory that in this case was named `_lus_scratch_ripple_mkdb_sp2b_25k_DAY_MONTH_DATE_TIME_ZONE_YEAR`. The shutdown would have removed all processes started by the Java execution, as could be verified by running `ps aux |grep $USER` from a command line after the "exiting..." message appeared.

## 12.5 Limitations of CGE Java API

### Network access requirements

Connecting to the database requires network access from the machine running the API to the node where the database is running. If there is no such access, no operations can be carried out.

### Launching the database

Launching a database relies upon being able to use the `cge-launch` command, this imposes two key limitations:

1. The command must be visible on the user's `$PATH` or the `$PATH` of the execution environment in order to launch a database, where `$PATH` refers to the `$PATH` environment variable.
2. Launching a database can only be carried out on the system i.e., if a user is running code that uses the API on a remote system, the user will not be able to launch databases.

## Accessing query results

When executing queries, the database writes the results to the configured file system. In order to retrieve those results from the API, access is required to the same file system and sufficient privileges are required to read those result files.

If queries are being executed on a machine without access to the configured file system, the user will only be able to access meta data about the results, not the results themselves.

## Log capture

The log capture functionality of the API relies upon access to the database log file. If that file is not known or not accessible (for example if the user is running on a remote machine), it will not be possible to retrieve log entries.

Note that the API will inject a unique identifier into the logs for each operation carried out via the API. Therefore it is possible to extract the log entries at a later date because the API will still be able to indicate the unique identifier used, which can be stored for later reference.

## 13 Logging and Troubleshooting

CGE produces a text log, which is a trace of program execution during query or update processing. Users can view the log with a text editor (such as vi), or typically the Linux `less` command. The log can be searched using the `grep` command for text messages of interest.

INFO messages will be deposited into the log during normal operation. CGE can also generate `ERROR` and `WARN` messages. All of these messages can yield information about activity that takes place during command execution.

System error message can be present in the log under conditions where CGE exits or improperly shuts down.

When queries or updates are executed, INFO messages with “now starting query #” are written to the log. For example:

```
2015-Feb-10 19:34:26.513 CST INFO [] [7720] 0x43 parser/parseAndBuildSM.cpp@374 allocQueryGlobals [] [QRY ] <OT> now starting query # 1
```

Many other INFO messages will also be deposited to the log during normal operation. For example, long processing times can be seen in the log from one INFO message to the next:

```
2015-Feb-13 14:44:45.500 CST INFO [] [9448] 0xb utils/malloc/cqe_malloc.cpp@901 LogRequest [] [QRY |MEM ] image 0 : request by "file: parser/qengine/database.cpp, func: readFromDisk line: 989" of 69.849 MiB (0x45d9688) was filled. (0x10005200c80)
2015-Feb-13 14:49:31.099 CST INFO [] [9448] 0xc parser/qengine/database.cpp@1141 readFromDisk [] [QRY |STRT] time to read in db of size 139.698 GiB (0x22ecb28000): 285.679279
```

When large datasets are used, the INFO message for the total start up time can be long, as shown in the following example:

```
2014-Dec-18 14:40:37.428 CST INFO [] [25977] 0x5b parser/dbServer.cpp@1259 main [] [QRY |STRT|PERF] Total startup time: 1434.489315 seconds
```

The following are examples of `ERROR` messages that CGE can produce when query or update processing has failed:

1. No such file or directory
2. No space left on device
3. Exiting because malloc of
4. Lookup failure for HURI
5. Invalid graph algorithm name
6. Exiting with status
7. Bad entry
8. Short read
9. Assertion
10. Realloc of
11. Error detected in Dispatcher

It is recommend to search the log for the text: “ERROR” and contact Cray Support if problems are encountered in query or update processing.

The following are samples of `WARN` messages that can be produced. `WARN` messages are subjective in preceding errors in processing:



1. huri was not found
2. directory not specified
3. not found in IRA
4. No valid quads in database
5. Invalid object for quad
6. Number of warnings found
7. Unsupported datatype
8. not in the dictionary
9. IRA huris not allocated

Search the log for `WARN` messages and contact Cray Support if problems in query or update processing are suspected.

The following are examples of system error messages that CGE can produce when query or update processing has failed. Search the log for the last `INFO` messages and contact Cray Support if any of these follow:

1. DUE TO TIME LIMIT
2. terminate called without an active exception
3. srun: error
4. Segmentation fault
5. Bus error
6. free invalid pointer
7. Out of memory
8. Unable to terminate gracefully
9. Floating point exception
10. Aborted
11. Killed
12. Unable to allocate resources
13. Exited with exit code
14. Requested nodes are busy
15. transaction completed with an error state
16. LIBDMAPP ERROR
17. IRI Resolution Error
18. rpn not found for
19. Trapped with SIGINT

## 13.1 Troubleshooting Common Cray Graph Engine (CGE) Issues

The most common errors that are likely to be encountered while using CGE involve failure to connect to a database server successfully. There are a variety of different errors that can occur, depending on exactly what

goes wrong. Common error messages that are likely to be encountered along with troubleshooting techniques are documented in the following table.

*Table 15. CGE Error Messages and Troubleshooting Information*

Error Message	Description	Resolution
Unable to establish a connection to the database server at host:port as it does not appear to be running	<p>The CLI tried to connect to a database server running on the given host and port combination but was unable to establish a connection. This typically means one of two things:</p> <ol style="list-style-type: none"> <li>1. There is no database server running on that host and port</li> <li>2. Firewall rules are preventing access to that host and port</li> </ol>	<ul style="list-style-type: none"> <li>• Verify that you have passed the correct host and port to the CLI</li> <li>• Verify that there is a database server running on that host and port</li> <li>• Verify that there are no firewall rules that are preventing access to the host and port. Contact a system administrator for additional information.</li> </ul>
Unable to authenticate to the database server at host:port. You do not have any SSH keys present in your configured identity idDirectory	The CLI tried to connect to a database server running on the given host and port combination. A connection was established successfully, but authentication to the database server failed because there are no SSH keys configured.	Create at least one SSH key and place it in the appropriate directory.
Unable to authenticate to the database server at host:port. Your SSH key(s) from your configured identity directory are not in the authorized_keys file of the database or its owner	<p>The CLI tried to connect to a database server running on the given host and port combination. A connection was established successfully but authentication to the database server failed because none of the SSH keys were in the <code>authorized_keys</code> file that the database is using.</p> <p>This may also be caused by the CLI selecting the wrong SSH identity. As described in the SSH identities section, the first identity found by searching several default locations is used, but this may not always be the desired identity.</p>	<ul style="list-style-type: none"> <li>• Review the database logs (if possible) to see which <code>authorized_keys</code> file was in-use: <ul style="list-style-type: none"> <li>◦ If the database server was launched, then this is either in the database directory itself or in the <code>~/ .cge</code> directory</li> <li>◦ If another user launched the database server, contact them to find out which <code>authorized_keys</code> file is in-use</li> </ul> </li> <li>• Add the public key to the relevant <code>authorized_keys</code> file, or ask the relevant user to do so.</li> </ul>

Error Message	Description	Resolution
		<ul style="list-style-type: none"> <li>Use the <code>--identity</code> option to specify the desired identity directory to use</li> </ul>
Host key for host host:port is not trusted, please run in interactive mode and trust this key or manually add the host key to your known_hosts file in your configured identity idDirectory	<p>The CLI tried to connect to a database server running on the given host and port combination. A connection was successfully established but the database server was unable to prove its identity to the CLI because the host key provided by the database server was not trusted.</p> <p>This error is usually only seen the first time when a connection to a specific server instance is established. Once the key is trusted (see resolution steps) this error should no longer be seen for this host and port combination.</p>	<ul style="list-style-type: none"> <li>If CGE is being run in interactive mode, the system will prompt to trust the host key. Enter <code>Yes</code> to do so.</li> <li>If it is required to use CGE non-interactively, adding the <code>--trust-keys</code> option to commands will automatically trust previously unknown host keys</li> </ul>
Timed out attempting to establish a database connection (waited N seconds), database server may be too busy to service your request currently	<p>The CLI tried to connect to a database server running on the given host and port combination but was unable to establish a connection within the timeout interval. This means that the database server is currently busy processing another request and cannot accept the request at this time.</p>	<ul style="list-style-type: none"> <li>Check the database logs to see what the database is currently doing <ul style="list-style-type: none"> <li>If the last log message states: "Trying to read RPN message from network..." then the database is ready, otherwise the database is busy</li> </ul> </li> <li>If the database is busy, there are a number of options that can be used to troubleshoot the issue: <ul style="list-style-type: none"> <li>Execute the request again later</li> <li>Increase the timeout with the <code>--timeout</code> option to wait for a longer period of time.</li> <li>Disable the timeout by setting <code>--timeout 0</code> to</li> </ul> </li> </ul>

Error Message	Description	Resolution
		<p>wait indefinitely until the database server is ready to process the next request</p> <ul style="list-style-type: none"> <li>In rare cases, the database may have become hung (if it is busy and you have not see any new log messages for long periods of time then this is most likely the problem) in which case you should kill and restart the database server and then retry your commands</li> </ul>
Server failed to start up	One or more of the CGE job steps failed to launch because CGE was not found.	<p>Try relaunching CGE if the system displays this message. In addition, it is recommended to ensure that all compute nodes are correctly configured. In particular verify the following:</p> <ul style="list-style-type: none"> <li>The same version of CGE is installed on all compute nodes and the login nodes</li> <li>All shared file systems are mounted and mounted in the same place on all compute nodes and the login nodes</li> <li>The munged process is running on all compute nodes</li> </ul> <p>If any of the preceding is not true and if relaunching the CGE CLI does not correct the problem, contact Cray Support.</p>
Not enough symmetric heap for new sorting keys	There is not enough symmetric heap for new sorting keys	<p>use the <code>-H</code> option to <code>cge-launch</code> to set the symmetric heap value to a larger value. Try doubling what shows up by default near the top of the log for a start.</p> <p>Symmetric heap is a boundary value on a resource that is allocated as needed, so using a larger than necessary value does not mean that this value will be allocated. It only means that no more than this value will be allocated. It is better to overestimate by a bit than to underestimate.</p>

Error Message	Description	Resolution
[PE_64]:inet_listen_socket_setup:inet_setup_listen_socket: bind failed port 20219 listen_sock = 5 Address already in use	This may be due to leftover <code>cge-server</code> processes	Follow the instructions documented in <a href="#">Terminate Orphaned cge-server Jobs</a> on page 119
Error: Timed out waiting for the server to start running	When a computational loop during a database build takes an extremely long time without producing any indication of forward progress (generally some kind of output in the log), <code>cge-launch</code> may decide that the start up sequence has hung and terminate it with this message.	Change the interval used to detect a start up hang from its default setting of 300 seconds (5 minutes) to some longer interval. If you know the problem is just that a dataset is very computationally intensive to build and is prone to such timeouts, setting this timeout value to 3600 seconds (an hour) is almost certain to eliminate any chance of this failure at the expense of causing you to take a very long time to detect an actual hang in start up. To change this, use the <code>--startupTimeout=seconds</code> option to <code>cge-launch</code> .
HTTP Errors are reported by a tool or API	A request submitted to the HTTP Interface provided by the <code>cge-cli fe</code> command was not successful. If the request was submitted via a tool or API then only minimal error details may be reported directly to you. However please see the resolutions for ways to find more detailed error information.	<ul style="list-style-type: none"> <li>• Submit the same request using a browser. The browser window may contain additional error messages which indicate the underlying error. Please review these carefully since they may indicate one of the other common errors detailed in this table.</li> <li>• Please review the front end logging as this will have logged the HTTP error and associated error details. These may indicate one of the other common errors detailed in this table.</li> <li>• If there is no obvious cause or additional error messages in the browser/front end logs then please review the database logs for error messages that may indicate if/ why the request failed on the database server.</li> <li>• In rare cases, the offending request may have caused the database server to crash in</li> </ul>

Error Message	Description	Resolution
		<p>which case, it will be necessary to relaunch it before making further requests</p> <ul style="list-style-type: none"> <li>○ If a crash has occurred please report this to your Cray support representative</li> </ul>
<pre>:inet_listen_socket_setup :inet_ setup_listen_socket : bind failed port 1371 listen_sock = 5 Address already in use</pre>	<p>A previous <code>cge-launch</code> or <code>HPC/mrun</code> job failed or was killed, and the <code>inet_listen_socket</code> is likely in the <code>TIME_WAIT</code> state on one or more of the compute nodes.</p>	<p>Wait 60-90 seconds for the <code>inet_listen_socket</code> (port 1371) to clear up from <code>TIME_WAIT</code> state. If the problem persists, the likely cause is some other program has an active socket connection to port 1371 on one (or more) compute nodes. That application must release port 1371 on the affected node(s) before new <code>cge-launch</code> or <code>HPC/mrun</code> jobs can be run on that node(s).</p>
<pre>2016-Jul-26 08:50:23.949 CDT WARN [][30526] 0x4e2 dictionary/ cqe_literals_info.cc@935 get_expr_result [] [QRY  DICT] unexpected type 31 in get_expr_result</pre>	<p>This may be caused by the <code>YEAR()</code> function when a user enters literals with two dashes for the date ("<code>--</code>"<sup>^</sup><code>xsd:date</code>).</p>	<p>Change the log level ad shown below: #0="leave at default", 1=Off (no output), 2=Error, 3=Warn, 4=Info, 5=Debug, 6=Trace</p> <pre>cge.server.DefaultLogLevel=4 cge.server.SetErrorKeywordsOn=0 x7FFFFFFFFFFFFFFF # Turn OFF QRY and DICT WARNings cge.server.SetWarnKeywordsOn= 0x7FFFFFFFFFDFFFE cge.server.SetInfoKeywordsOn=0x 7FFFFFFFFFFFFFFF</pre>

If the same error is encountered even after following the suggested fixes in this section it is recommended to add the `--trace` option to the command in order to get detailed information about the communications being attempted and review the log messages carefully both on the front end and server side to try and understand what is going wrong. It is also worth reviewing the server logs if you are able as there are some situations, which will manifest as client side communications errors caused by an error on the serve. Reviewing the server logs may provide additional information about why you are encountering an error. If you are still unable to resolve the issue, please contact Cray support providing logs from both the CLI and the database server to aid in diagnosis of the issue.

The results format received in the browser is dictated by the HTTP Accept header that your browser sends (or conversely that your programmatic HTTP client sends). The `Force text/plain` as the response `Content` option controls the `Content-Type` header that the front end responds with, which affects how the browser interprets the response. Depending on the browser if this option is disabled (the default) then this might mean that it downloads/offers to save the response to a file rather than displaying it in the browser, enabling the aforementioned option changes the response `Content-Type` to always be `text/plain` regardless of what format the front end actually outputs which forces the browser to display the response in the browser itself. If the results

need to displayed in a different format, customize the HTTP Accept header accordingly. Most browsers have some means to configure this. For example, in Firefox, navigate to **About>Config**. Click through the warning if it appears, and then search for accept and edit the value of the `network.http.accept.default` setting to add the desired content types. The closest thing to plain text that the front end will produce is text/tab-separated-values. Most browsers include `application/xml` in their default accept header, which means you will typically get SPARQL XML results by default (or RDF/XML if it were a `CONSTRUCT` query).

## Additional Important Information

- If the same error is encountered even after following the suggested fixes in this section it is recommended to add the `--trace` option to the command in order to get detailed information about the communications being attempted and review the log messages carefully both on the front end and server side to try and understand what is going wrong. It is also worth reviewing the server logs if you are able as there are some situations, which will manifest as client side communications errors caused by an error on the server. Reviewing the server logs may provide additional information about why you are encountering an error. If you are still unable to resolve the issue, please contact Cray support providing logs from both the CLI and the database server to aid in diagnosis of the issue.
- The results format received in the browser is dictated by the HTTP Accept header that your browser sends (or conversely that your programmatic HTTP client sends). The `Force text/plain as the response Content` option controls the `Content-Type` header that the front end responds with, which affects how the browser interprets the response. Depending on the browser if this option is disabled (the default) then this might mean that it downloads/offers to save the response to a file rather than displaying it in the browser, enabling the aforementioned option changes the response `Content-Type` to always be text/plain regardless of what format the front end actually outputs which forces the browser to display the response in the browser itself. If the results need to displayed in a different format, customize the HTTP Accept header accordingly. Most browsers have some means to configure this. For example, in Firefox, navigate to **About>Config**. Click through the warning if it appears, and then search for accept and edit the value of the `network.http.accept.default` setting to add the desired content types. The closest thing to plain text that the front end will produce is text/tab-separated-values. Most browsers include `application/xml` in their default accept header, which means you will typically get SPARQL XML results by default (or RDF/XML if it were a `CONSTRUCT` query).
- Writing a large dataset to a Lustre directory when there is insufficient free space on the Lustre file system may produce a segmentation fault. To resolve this issue, free up some more space on the file system.
- When an input triple or quad data file does not end with a newline character, the system will return an issue. Therefore, it is important to ensure that the `.nt/.nq` file ends with a newline.

## 13.2 Terminate Orphaned cge-server Jobs

### Prerequisites

This procedure requires root privileges.

### About this task

Follow the instructions listed in this procedure to track orphaned `cge-server` jobs down and terminate them. The examples shown in this procedure can be used for a system with 3 sub-racks.

## Procedure

1. Log on to the System Management Workstation (SMW) as root

2. Execute the following to find out if there are stray `cge-server` processes.

```
# pdsh -w 'nid000[00-47]' "ps -ef|grep 'cge-serve[r]'|grep -v grep | awk '{print \$2}';true"|wc -w
```

3. Terminate the stray `cge-server` processes

```
# pdsh -w nid000[00-30,32-46] "ps -ef|grep 'cge-serve[r]'|awk '{print \$2}'|xargs kill"
```

4. Rerun the preceding command to ensure all stray `cge-server` processes have been terminated.

5. Verify that all the stray `cge-server` processes have been terminated by executing the following command:

```
# pdsh -w 'nid000[00-47]' "ps -ef|grep 'cge-serve[r]'|grep -v grep | awk '{print \$2}';true"|wc -w
0
```

This output indicates that everything has been cleared.