

IBM Business Analytics

IBM Cognos Analytics 11.1.7 Performance Tuning and Monitoring

Performance Team
Cognos Analytics 11.1.7

Cognos And System Performance



- All customers desire the capability to squeeze the maximum performance out of their IBM Cognos Analytics (CA) investment.
- IBM Cognos Analytics is a complex software and hardware environment. A single bottleneck in either the software or hardware has a ripple effect for the entire system.
- We have compiled a list of the most common software and hardware bottlenecks and will go over how to monitor the system for them.
 - Process Threading:
 - Too few threads and queuing can occur.
 - Process Memory Sizing:
 - Too little memory and OutOfMemory conditions can exist.
 - Too little memory and Garbage Collection can become costly.
- Nothing can minimize the impact of a sub-optimal Analytics model or poorly authored report spec.

Outline / Agenda

- Cognos Analytics Architecture
- Cognos Analytics Performance
 - Behind the Scenes
 - End User View
- JVM Tuning Comparison
- System Performance
- Monitoring Cognos Analytics

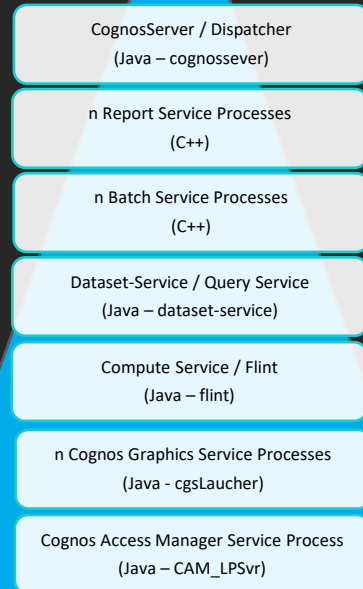


Cognos Analytics Architecture

Behind the Scenes

Cognos Analytics Basic Architecture

The following Services account for the core Analytics processes seen in system process views and are generally tunable.



CognosServer / Dispatcher – Java application responsible for routing requests through the Analytics system and managing Analytics content (cognossever).

Report Service – C++ application that manages interactive requests to execute reports (BIBusTKServerMain).

Batch Report Service – C++ application that manages non-interactive requests to execute reports (BIBusTKServerMain).

Dataset-Service / Query Service– Java application that manages Query Service requests and returns the result to the requesting Service (dataset-service).

Compute Service / Flint – Child java application to the Query Service used for parquet related queries (flint)

Cognos Graphics Service – Java application that produces graphics on behalf of the other services (cgsLauncher).

Cognos Access Manager Service – Java application that handles user authentication, authorization, and encryption (CAM_LPSvr).

In addition to the above services there are a number of micro services embedded within the java processes.

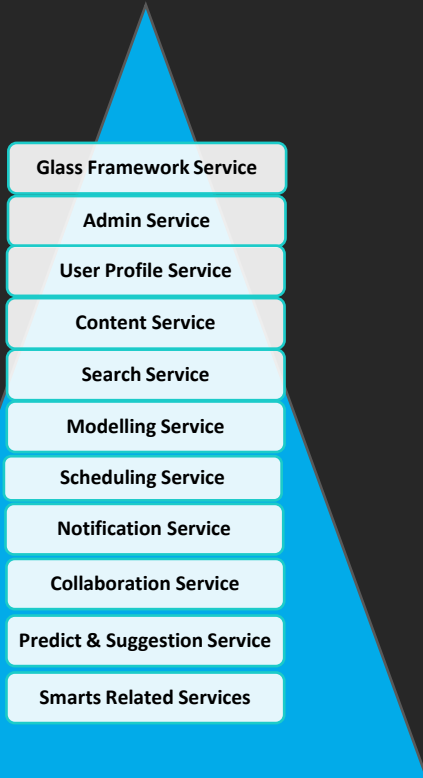
Cognos Analytics Micro-Services

The following are some of the micro-services available in Cognos Analytics 11.1.7. They are embedded in the existing java processes and are not necessarily readily exposed for tuning.

- **Glass Framework Service**
- **Admin Service**
- **User Profile Service**
- **Content Service**
- **Search Service**
- **Modelling Service**
- **Scheduling Service**
- **Notification Service**
- **Collaboration Service**
- **Predict & Suggestion Service**
- **Many Smarts Related Services (Conversation, Explore, Insight, Knowledge & Discovery, Visualization & Recommender)**

All CA micro services run in the JVM processes and are installed with all “app-server” installations.

A full list of the microservices can be found in < Cognos root >/configuration/bi-services/*



Cognos Analytics Processes (Windows)



Cognos Analytics 11.1.7 system under load (Windows):

cogbootstrapservice.exe	43698	< 0.01	108,032 K	31,090 K	31,116 K	339	5 "D:\IBM\cognos\c10_64_EndorR1_rs\bin64\cogbootstrapservice.exe"
java.exe	26852	17.61	20,878,240 K	7,749,856 K	7,853,244 K	14,109	1,881 ..\jre\bin\java.exe -Xcompressedrefs -Djava.net.preferIPv4Stack=true -Xmx16384m -Xms8192m -
conhost.exe	37836	< 0.01	2,147,509,796 K	3,288 K	3,308 K	51	2 \??\C:\Windows\system32\conhost.exe 0x4
cgsLauncher.exe	32652		17,920,084 K	1,044,308 K	1,264,964 K	785	53 D:\IBM\cognos\c10_64_EndorR1_rs\bin64\cgsLauncher.exe -option COG_ROOT=D:\IBM\cog
conhost.exe	48620		2,147,509,496 K	2,940 K	2,960 K	49	1 \??\C:\Windows\system32\conhost.exe 0x4
java.exe	3884	13.49	22,584,652 K	3,704,068 K	3,874,380 K	2,648	276 D:\IBM\cognos\c10_64_EndorR1_rs\jre\bin\java.exe -Xms20480m -Xmx20480m -Xmx6826m -
conhost.exe	34092	< 0.01	2,147,509,496 K	2,964 K	2,984 K	49	1 \??\C:\Windows\system32\conhost.exe 0x4
java.exe	7352	< 0.01	9,273,536 K	381,200 K	487,724 K	968	120 D:\IBM\cognos\c10_64_EndorR1_rs\jre\bin\java.exe -Xms1024m -Xmx8192m -XX:MaxDirectMe
conhost.exe	37720		2,147,509,496 K	2,972 K	2,992 K	49	1 \??\C:\Windows\system32\conhost.exe 0x4
BIBusTKServerMain.exe	4320	< 0.01	644,820 K	301,556 K	301,556 K	747	46 D:\IBM\cognos\c10_64_EndorR1_rs\bin\BIBusTKServerMain.exe threads=20 camssl=false C0C
conhost.exe	44576		2,147,509,500 K	2,936 K	2,956 K	49	1 \??\C:\Windows\system32\conhost.exe 0x4
BmtMDProviderMain.exe	23260	< 0.01	145,340 K	28,728 K	29,672 K	361	13 D:\IBM\cognos\c10_64_EndorR1_rs\bin\BmtMDProviderMain.exe threads=10 camssl=false C0C
conhost.exe	41556		2,147,509,500 K	2,932 K	2,952 K	49	1 \??\C:\Windows\system32\conhost.exe 0x4
BIBusTKServerMain.exe	1280	1.24	562,744 K	218,204 K	218,204 K	690	46 D:\IBM\cognos\c10_64_EndorR1_rs\bin\BIBusTKServerMain.exe threads=20 camssl=false C0C
conhost.exe	31268		2,147,509,500 K	2,924 K	2,924 K	49	1 \??\C:\Windows\system32\conhost.exe 0x4
BIBusTKServerMain.exe	9724	1.25	89,448 K	21,788 K	21,788 K	293	3 D:\IBM\cognos\c10_64_EndorR1_rs\bin\BIBusTKServerMain.exe threads=20 camssl=false C0C
conhost.exe	8848		2,147,509,500 K	2,924 K	2,924 K	50	1 \??\C:\Windows\system32\conhost.exe 0x4

Cognos Analytics Processes (Linux/Unix)

Cognos Analytics 11.1.7 system under load (AIX):

```
root@phantaia:~# ps -eo pid,ppid,cmd,stat,cpu,mem --sort=-cpu | head -n 1000
  PID   PPID  CMD                               STAT  CPU   MEM
  570224  65.4  3733M 3733M 128 3733M63878 1% 0 126 0 java
  7995796 326.5 3580M 3580M 128 3580M20373 1% 0 4 0 java
  8257862 12.8 994M 994M 64 994M 1 0% 0 481 0 IBusTKServerMain
  5833296 0.0 478M 478M 128 477M 0 0% 0 0 0 java
  7078510 12.9 273M 273M 64 273M 1 0% 0 585 0 IBusTKServerMain
  8323466 13.7 271M 271M 64 271M 1 0% 0 514 0 IBusTKServerMain
  6357798 17.3 268M 268M 64 268M 1 0% 0 790 0 IBusTKServerMain
  7143772 12.1 260M 260M 64 260M 1 0% 0 449 0 IBusTKServerMain
  8127092 12.3 259M 259M 64 259M 1 0% 0 510 0 IBusTKServerMain
  6685604 20.7 258M 258M 64 258M 0 0% 0 726 0 IBusTKServerMain
  8651194 12.8 258M 258M 64 258M 1 0% 0 547 0 IBusTKServerMain
  8454484 14.2 255M 255M 64 255M 1 0% 0 561 0 IBusTKServerMain
  5768124 14.2 255M 255M 64 255M 1 0% 0 573 0 IBusTKServerMain
  8389632 17.2 253M 253M 64 253M 1 0% 0 648 0 IBusTKServerMain
  9045148 12.5 252M 252M 64 252M 1 0% 0 510 0 IBusTKServerMain
  8782934 12.8 249M 249M 64 249M 1 0% 0 513 0 IBusTKServerMain
  8192278 18.6 247M 247M 64 247M 1 0% 0 903 0 IBusTKServerMain
  6161202 20.3 244M 244M 64 244M 1 0% 0 837 0 IBusTKServerMain
  5309944 12.5 243M 243M 64 243M 1 0% 0 695 0 IBusTKServerMain
  8061398 56.9 144M 144M 128 144M 778 0% 0 15 0 java
  7209376 55.9 138M 138M 128 137M 836 0% 0 0 0 java
  7733580 54.4 137M 137M 128 137M 810 0% 0 4 0 java
  7930158 68.8 133M 133M 128 133M 981 0% 0 490 0 java
  2032170 0.0 56716 54556 4388 50168 0 0% 0 0 0 STAFProc
  1311954 0.0 25216 25216 0 25216 0 0% 0 0 0 j2pg
  7405850 0.0 17284 17648 628 17020 0 0% 0 0 0 cogbootstrapservice
  7341066 0.0 14560 14644 124 14520 0 0% 0 0 0 BmtMDPProviderMain
```

```

  7995796 326.5 3580M 3580M 128 3580M20373 1% 0 4 0 java
  8257862 12.8 994M 994M 64 994M 1 0% 0 481 0 IBusTKServerMain
  5833296 0.0 478M 478M 128 477M 0 0% 0 0 0 java
  7078510 12.9 273M 273M 64 273M 1 0% 0 585 0 IBusTKServerMain
  8323466 13.7 271M 271M 64 271M 1 0% 0 514 0 IBusTKServerMain
  6357798 17.3 268M 268M 64 268M 1 0% 0 790 0 IBusTKServerMain
  7143772 12.1 260M 260M 64 260M 1 0% 0 449 0 IBusTKServerMain
  8127092 12.3 259M 259M 64 259M 1 0% 0 510 0 IBusTKServerMain
  6685604 20.7 258M 258M 64 258M 0 0% 0 726 0 IBusTKServerMain
  8651194 12.8 258M 258M 64 258M 1 0% 0 547 0 IBusTKServerMain
  8454484 14.2 255M 255M 64 255M 1 0% 0 561 0 IBusTKServerMain
  5768124 14.2 255M 255M 64 255M 1 0% 0 573 0 IBusTKServerMain
  8389632 17.2 253M 253M 64 253M 1 0% 0 648 0 IBusTKServerMain
  9045148 12.5 252M 252M 64 252M 1 0% 0 510 0 IBusTKServerMain
  8782934 12.8 249M 249M 64 249M 1 0% 0 513 0 IBusTKServerMain
  8192278 18.6 247M 247M 64 247M 1 0% 0 903 0 IBusTKServerMain
  6161202 20.3 244M 244M 64 244M 1 0% 0 837 0 IBusTKServerMain
  5309944 12.5 243M 243M 64 243M 1 0% 0 695 0 IBusTKServerMain
  8061398 56.9 144M 144M 128 144M 778 0% 0 15 0 java
  7209376 55.9 138M 138M 128 137M 836 0% 0 0 0 java
  7733580 54.4 137M 137M 128 137M 810 0% 0 4 0 java
  7930158 68.8 133M 133M 128 133M 981 0% 0 490 0 java
  2032170 0.0 56716 54556 4388 50168 0 0% 0 0 0 STAFProc
  1311954 0.0 25216 25216 0 25216 0 0% 0 0 0 j2pg
  7405850 0.0 17284 17648 628 17020 0 0% 0 0 0 cogbootstrapservice
  7341066 0.0 14560 14644 124 14520 0 0% 0 0 0 BmtMDPProviderMain


  # ps -eo pid,ppid,cmd,stat,cpu,mem --sort=-cpu | tail -n 1000
  7995796 326.5 3580M 3580M 128 3580M20373 1% 0 4 0 java
  8257862 12.8 994M 994M 64 994M 1 0% 0 481 0 IBusTKServerMain
  5833296 0.0 478M 478M 128 477M 0 0% 0 0 0 java
  7078510 12.9 273M 273M 64 273M 1 0% 0 585 0 IBusTKServerMain
  8323466 13.7 271M 271M 64 271M 1 0% 0 514 0 IBusTKServerMain
  6357798 17.3 268M 268M 64 268M 1 0% 0 790 0 IBusTKServerMain
  7143772 12.1 260M 260M 64 260M 1 0% 0 449 0 IBusTKServerMain
  8127092 12.3 259M 259M 64 259M 1 0% 0 510 0 IBusTKServerMain
  6685604 20.7 258M 258M 64 258M 0 0% 0 726 0 IBusTKServerMain
  8651194 12.8 258M 258M 64 258M 1 0% 0 547 0 IBusTKServerMain
  8454484 14.2 255M 255M 64 255M 1 0% 0 561 0 IBusTKServerMain
  5768124 14.2 255M 255M 64 255M 1 0% 0 573 0 IBusTKServerMain
  8389632 17.2 253M 253M 64 253M 1 0% 0 648 0 IBusTKServerMain
  9045148 12.5 252M 252M 64 252M 1 0% 0 510 0 IBusTKServerMain
  8782934 12.8 249M 249M 64 249M 1 0% 0 513 0 IBusTKServerMain
  8192278 18.6 247M 247M 64 247M 1 0% 0 903 0 IBusTKServerMain
  6161202 20.3 244M 244M 64 244M 1 0% 0 837 0 IBusTKServerMain
  5309944 12.5 243M 243M 64 243M 1 0% 0 695 0 IBusTKServerMain
  8061398 56.9 144M 144M 128 144M 778 0% 0 15 0 java
  7209376 55.9 138M 138M 128 137M 836 0% 0 0 0 java
  7733580 54.4 137M 137M 128 137M 810 0% 0 4 0 java
  7930158 68.8 133M 133M 128 133M 981 0% 0 490 0 java
  2032170 0.0 56716 54556 4388 50168 0 0% 0 0 0 STAFProc
  1311954 0.0 25216 25216 0 25216 0 0% 0 0 0 j2pg
  7405850 0.0 17284 17648 628 17020 0 0% 0 0 0 cogbootstrapservice
  7341066 0.0 14560 14644 124 14520 0 0% 0 0 0 BmtMDPProviderMain

  # ps -eo pid,ppid,cmd,stat,cpu,mem --sort=-cpu | head -n 1000
  7995796 326.5 3580M 3580M 128 3580M20373 1% 0 4 0 java
  8257862 12.8 994M 994M 64 994M 1 0% 0 481 0 IBusTKServerMain
  5833296 0.0 478M 478M 128 477M 0 0% 0 0 0 java
  7078510 12.9 273M 273M 64 273M 1 0% 0 585 0 IBusTKServerMain
  8323466 13.7 271M 271M 64 271M 1 0% 0 514 0 IBusTKServerMain
  6357798 17.3 268M 268M 64 268M 1 0% 0 790 0 IBusTKServerMain
  7143772 12.1 260M 260M 64 260M 1 0% 0 449 0 IBusTKServerMain
  8127092 12.3 259M 259M 64 259M 1 0% 0 510 0 IBusTKServerMain
  6685604 20.7 258M 258M 64 258M 0 0% 0 726 0 IBusTKServerMain
  8651194 12.8 258M 258M 64 258M 1 0% 0 547 0 IBusTKServerMain
  8454484 14.2 255M 255M 64 255M 1 0% 0 561 0 IBusTKServerMain
  5768124 14.2 255M 255M 64 255M 1 0% 0 573 0 IBusTKServerMain
  8389632 17.2 253M 253M 64 253M 1 0% 0 648 0 IBusTKServerMain
  9045148 12.5 252M 252M 64 252M 1 0% 0 510 0 IBusTKServerMain
  8782934 12.8 249M 249M 64 249M 1 0% 0 513 0 IBusTKServerMain
  8192278 18.6 247M 247M 64 247M 1 0% 0 903 0 IBusTKServerMain
  6161202 20.3 244M 244M 64 244M 1 0% 0 837 0 IBusTKServerMain
  5309944 12.5 243M 243M 64 243M 1 0% 0 695 0 IBusTKServerMain
  8061398 56.9 144M 144M 128 144M 778 0% 0 15 0 java
  7209376 55.9 138M 138M 128 137M 836 0% 0 0 0 java
  7733580 54.4 137M 137M 128 137M 810 0% 0 4 0 java
  7930158 68.8 133M 133M 128 133M 981 0% 0 490 0 java
  2032170 0.0 56716 54556 4388 50168 0 0% 0 0 0 STAFProc
  1311954 0.0 25216 25216 0 25216 0 0% 0 0 0 j2pg
  7405850 0.0 17284 17648 628 17020 0 0% 0 0 0 cogbootstrapservice
  7341066 0.0 14560 14644 124 14520 0 0% 0 0 0 BmtMDPProviderMain
```


Cognos Analytics Performance

Behind the Scenes

Sizing Cognos Analytics



Ensuring the Cognos Analytics environment is sized properly is essential to having a stable and performant environment. It's recommended to start with the default JVM settings then monitor the system for any additional tuning required.

CA 11.1.X:

https://www.ibm.com/support/knowledgecenter/en/SSEP7J_11.1.0/com.ibm.swg.ba.cognos.inst_cr_winux.doc/c_installs_verify_ram.html

Use the following memory settings as a starting point and adjust them based on the memory usage of your system.


- 2 GB for the base operating system and accompanying software, such as antivirus, backup and enterprise management software
- 8 GB for the Dispatcher JVM (Content Manager or Application Tier)
- 2 GB for Cognos Graphics Service JVM
- 8 GB for Query Service / Dataset Service JVM
- 2 GB per BIBus for report server processes
- 1 GB initial heap size for Compute service for Datasets / 8 GB Maximum (default settings)

Cognos Analytics Memory Overview



- There are 4 JVM's associated with Cognos Analytics. By default they use the IBM Java 8 SR6 FP15 included with the product at <code>< cognos root >/jre</code>. The default (and recommended starting values) for the JVM's are the following:
 - (WLP) cognosserver / Dispatcher JVM: 8GB
 - Needed on all Application servers and the Content Manager
 - (WLP) dataset-service / Query Service JVM: 8GB
 - Needed on all Application Servers
 - Flint / Compute Service JVM: 8GB
 - Needed on all Application Servers. A sub process of the Query Service JVM.
 - Cognos Graphics Service JVM: 1GB
- The Report Service and Batch Report Service processes will dynamically vary in size as necessary. We recommend allocating 2GB per process initially.
- For a starting point we recommend leaving at least 4GB of memory available for the base operating system and any other mandatory software (i.e. back up software, antivirus, enterprise tools).

WebSphere Liberty Profiles (WLP) for Analytics



- The WebSphere Liberty Profiles bundled in Cognos Analytics 11.1.7 FP2 is Liberty version 20.0.0.7.
 - Located in the CA installation at <code>< cognos root >/wlp</code>
- There are two WebSphere Liberty Profiles:
 - CognosServer (Dispatcher JVM): <code>< cognos root >/wlp/usr/servers/cognosserver</code>
 - Dataset-Service (Query Service JVM): <code>< cognos root >/wlp/usr/servers/dataset-service</code>
- Non-optimal settings for the JVM's can lead to:
 - OutOfMemory conditions
 - Frequent or long Garbage Collection pauses
 - 'timeout' or 'unresponsive' application
 - Overall poor performance

WLP for Cognos Analytics

- Port allocation for the CognosServer is defined in Cognos Configuration. By default the CognosServer port is 9300.
- Port allocation for the Dataset-Service is defined in Cognos Configuration. By default the port is 9301.
- Port allocation for the Compute Service can be configured in Cognos Configuration.


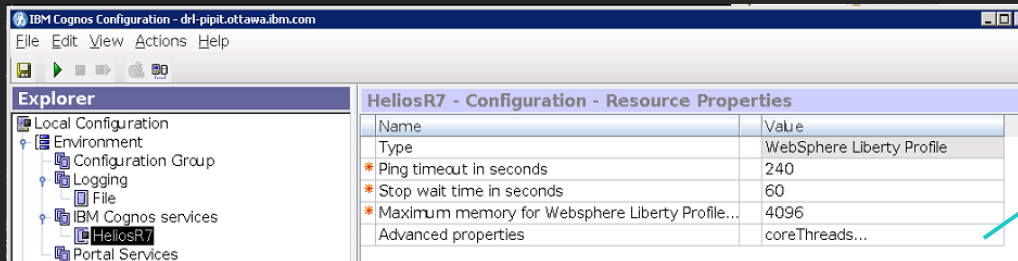
The screenshot shows the IBM Cognos Configuration Explorer interface. The left pane displays a tree view under 'Local Configuration' with 'Environment' selected. The right pane shows the 'Environment - Group Properties' table.

Name	Value
Gateway Settings	
* Gateway URI	https://dal-panther.ottawa.ibm.com:9443/ibmcognos/bi/v1/disp
Gateway namespace	
Content Manager sAMAccountName	
Allow namespace override?	False
* Dispatcher URIs for gateway	https://dal-phastos.ottawa.ibm.com:9300/bi/v1/disp
* Controller URI for gateway	https://dal-phastos.ottawa.ibm.com:80/ibmcognos/controllerServer
Dispatcher Settings	
* External dispatcher URI	https://dal-phastos.ottawa.ibm.com:9300/p2pd/servlet/dispatch
* Internal dispatcher URI	https://dal-phastos.ottawa.ibm.com:9300/p2pd/servlet/dispatch
Dispatcher password	*****
External JMX port	0
External JMX credential	*****
* Report Server execution mode	32-bit
Dataset Service Settings	
Dataset Service port number	9301
Compute Service port number	0
Other URI Settings	
* Dispatcher URI for external applica...	https://dal-phastos.ottawa.ibm.com:9300/bi/v1/disp
* Content Manager URIs	https://dal-panther.ottawa.ibm.com:9300/p2pd/servlet


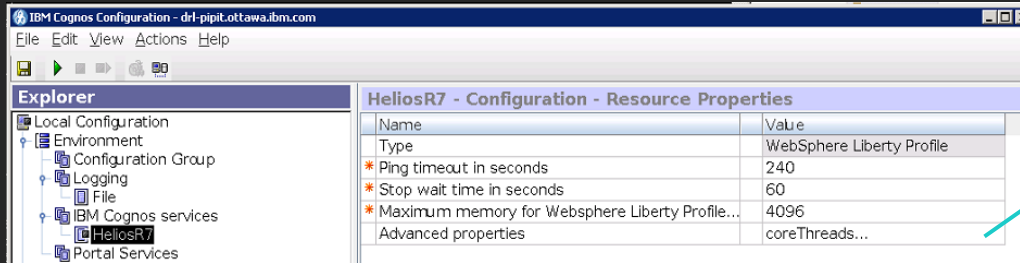
WLP for CA: Cognos Server



- The CognosServer JVM is configured in Cognos Configuration.
- The default WLP settings for CA 11 CognosServer (Dispatcher) java process:
 - JVM Settings
 - Initial Heap Size (Xms) = 1GB or 1024MB
 - Maximum Heap Size (Xmx) = 8GB or 8192MB
 - Xgcpolicy:gencon is set by default
 - <http://javaeesupportpatterns.blogspot.com/2012/03/ibm-jvm-tuning-gencon-gc-policy.html>
 - Xcompressedrefs is set by default (IBM Java in the <cognos root>/bin64/bootstrap_wlp_<os>.xml)
 - Thread Pools: coreThreads (100) / maxThreads (1500)




Cognos Server: Threads



- The property value of 'coreThreads' represents the number of threads that the WLP server starts up with.
- The 'maxThreads' value represents the maximum number of threads that can be associated with the WLP server.
- The setting of coreThreads = -1 means that at runtime, coreThreads is set to a multiple of the number of hardware threads on your system (<https://developer.ibm.com/wasdev/docs/was-liberty-threading-and-why-you-probably-dont-need-to-tune-it/>). More threads will be created as needed.
- If threads become exhausted messages will be written to the p2pd_messages.log on the application server.

[4/6/18 14:47:46:439 EDT] 0000005d com.ibm.ws.threading.internal.ThreadPoolController W CWWKE1200W: **All threads in the Liberty default executor appear to be hung. Liberty automatically increased the number of threads from 100 to 100. However, all threads still appear to be hung.**

WLP for CA: Dataset-Service



- The dataset-service JVM is configured in the Cognos Administration Console.
- The default WLP settings for the CA 11.1 dataset-service (Query Service) java process are:
 - JVM Settings
 - Initial Heap Size (Xms) = 1GB or 1024MB
 - Maximum Heap Size (Xmx) = 8GB or 8192MB
 - Xgcpolicy:gencon is set by default
 - <http://javaeesupportpatterns.blogspot.com/2012/03/ibm-jvm-tuning-gencon-gc-policy.html>
 - Xcompressedrefs is set by default

<input type="checkbox"/>	Category	Name	Value	Acquired
<input type="checkbox"/>	Tuning	Minimum query execution time before a result set is considered for caching (milliseconds)	50	Yes
<input type="checkbox"/>	Tuning	Initial JVM heap size for the query service (MB) (Requires QueryService restart)	1024	Yes
<input type="checkbox"/>	Tuning	JVM heap size limit for the query service (MB) (Requires QueryService restart)	8192	Yes
<input type="checkbox"/>	Tuning	Initial JVM nursery size (MB) (Requires QueryService restart)	0	Yes
<input type="checkbox"/>	Tuning	JVM nursery size limit (MB) (Requires QueryService restart)	0	Yes
<input type="checkbox"/>	Tuning	JVM garbage collection policy (Requires QueryService restart)	Generational	Yes
<input type="checkbox"/>	Tuning	Additional JVM arguments for the query service (Requires QueryService restart)		Yes
<input type="checkbox"/>	Tuning	Number of garbage collection cycles output to the verbose log (Requires QueryService restart)	1000	Yes
<input type="checkbox"/>	Tuning	Disable JVM verbose garbage collection logging (Requires QueryService restart)	<input type="checkbox"/>	Yes

Dataset-Service / Query Service



- Threading is handled dynamically by the Query Engine. The number of Interactive Report Service Processes and Batch Report Service Processes affect the thread count.
- The Query Service automatically provides logs (dq_verbosegc_<timestamp>.log) to help determine if your values are set correctly and allow for easy debugging.
 - Located in the <cognos root>/logs/XQE folder

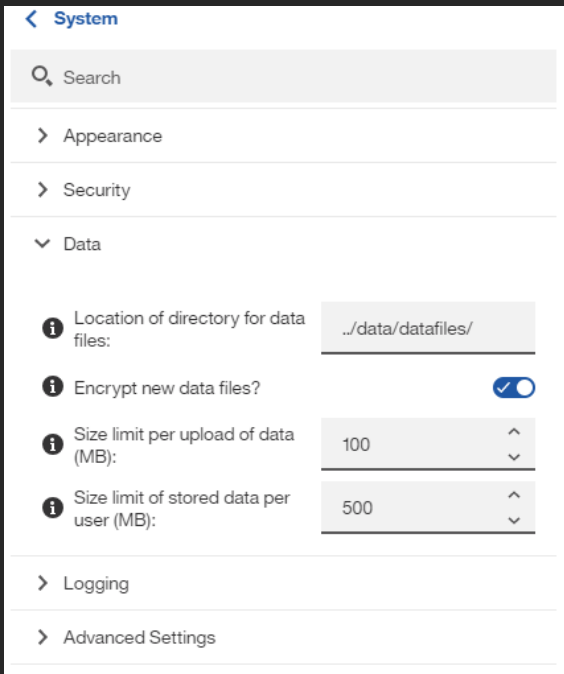
Dynamic Cubes

- Dynamic Cubes are an extension of the Query Engine that leverages substantial in-memory data assets as well as aggregate awareness in order to achieve high performance interactive analysis & reporting over terabytes of warehouse data.
- The JVM must be tuned correctly in order to obtain optimal performance depending on:
 - The size of the data it consumes
 - The expected workload



Compute Service / Flint

- The Compute Service JVM process is a sub-process launched and managed by the Query Service (dataset-service) JVM. The Compute Service JVM process is called flint. The Compute Service starts and stops when the Query Service starts and stops.
- The Compute Service gets called when Parquet related queries are issued in order to improve performance and caching. Queries using Parquet are stored locally on disk. The location of the files is configurable via Manage > Configuration > System > Data.
- The default JVM of 8GB and maximum threads of 1500 is configurable through the product files. Edit the xqe.config.xml file at <cognos root>/configuration:



```
<!-- The minimum and maximum amount of memory (MB) to be allocated to Flint Server.
Direct memory (off-heap) is used for Netty transport buffers. -->
<memory minHeap="1024" maxHeap="8192" maxDirect="512"/>
```

```
<!-- The amount of Thrift workers receiving requests. -->
<thriftWorkers min="5" max="1500"/>
```

Cognos Graphics Service (CGS) Tuning



- Threading for the Cognos Graphics Service is configured in the Cognos Administration Console: 50 threads per process.

Category	Name	Value	Acquired
Tuning	Number of low affinity connections for the graphics service during non-peak period	50	Yes
Tuning	Number of high affinity connections for the graphics service during peak period	1	Yes
Tuning	Number of low affinity connections for the graphics service during peak period	50	Yes

- The size of the JVM size is configured in a file on disk. By default it is 1GB in size with no JVM tuning applied.

Unix / Linux: Configured in the `cgsServer.sh` file in the `<cognos root>/bin` and `<cognos root>/bin64` locations. The JVM arguments are set after `$JAVA_OPTS`:

```
$JAVA_OPTS -Xmx2g -Xms2g -Xmn1g -Xcompressedrefs -Xgcpolicy:gencon
```

Windows: Configured in the `cgsService.xml` file in the `<cognos root>/webapps/p2pd/WEB-INF/services` location. The JVM arguments are set between `child-proc-cmd` tags after `vmargs` (2 places in the same file):

```
<child-proc-cmd>vmargs</child-proc-cmd>  
<child-proc-cmd>Xmx2g</child-proc-cmd>  
<child-proc-cmd>Xms2g</child-proc-cmd>  
<child-proc-cmd>Xmn1g</child-proc-cmd>  
<child-proc-cmd>Xcompressedrefs</child-proc-cmd>  
<child-proc-cmd>Xgcpolicy:gencon</child-proc-cmd>
```

Report Service Tuning



Key things to consider when tuning Report Service:

- Queuing can occur if there are not enough Report Service processes for the work load on the system.
- Queuing of requests negatively impacts the Analytics system.
 - Affects both Interactive Report Service and Batch Report Service.
 - Affects Dynamic and Compatible Reporting Engines.
- Monitoring via the Cognos Administration Console:

Metrics - System

0 0 0 No metric score

[-] Queue - Report service

Latency		00:00:00.015
Number of queue requests		802
Queue length	--	39
Queue length high watermark		39
Queue length low watermark		1
Time in queue		00:00:12.076
Time in queue high watermark		00:00:48.223


39 users in queue

Waiting 12s


Report Service Tuning



- The threading model in Cognos Analytics 11 Report Server defaults to 8 low affinity threads and 2 high affinity threads.
- With 2 Report Server processes there are 16 low affinity threads (default value).
- If 24 requests are issued to report service, 16 get served and 8 get queued.



- Increase the number of Report Server processes through the IBM Cognos Administration Console:



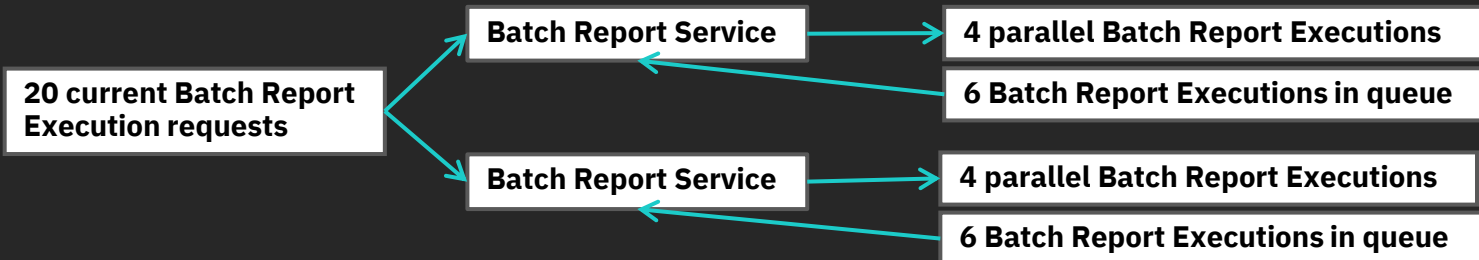
Solution – Increase the Report Service Processes to 3 to prevent Queuing!

Batch Report Service Tuning



- Similar to the Report Service. The Batch Report Service handles job report execution. The number of processes has a significant effect on batch report execution.
- Too few Batch Report Service processes will lead to report execution requests waiting in the Queue.
- Set in the IBM Cognos Administration Console:

Category	Name	Value	Acqu
Tuning	Maximum number of processes for the batch report service during non-peak period	2	Yes
Tuning	Maximum number of processes for the batch report service during peak period	2	Yes



Solution – Increase the Batch Report Service Processes to 5 to prevent Queuing!

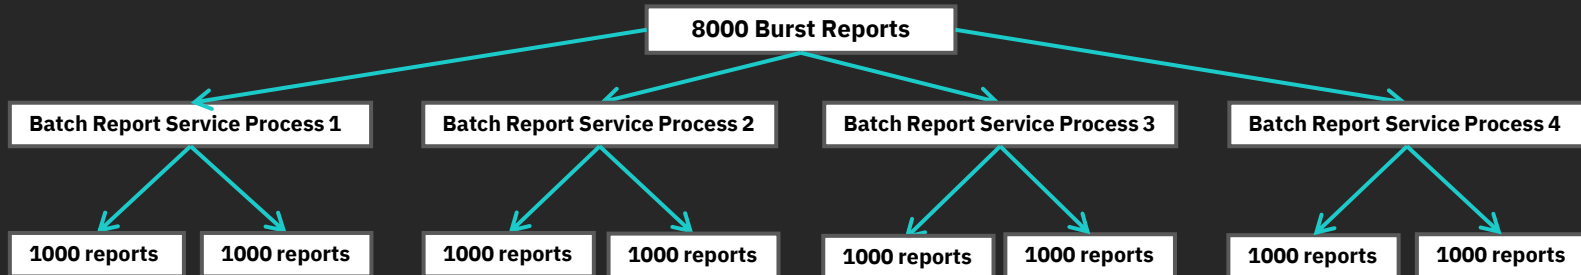
Batch Report Service Tuning



- The number of Delivery Service connections may need to be increased for heavy Batch environments that write to disk:

Category	Name	Value	Acqu
Tuning	Maximum connections for delivery service during non-peak period	4	Yes
Tuning	Maximum number of delivery service connections during peak period	4	Yes

- By default in CA11, Burst Report execution are as follows.
 - The work gets “chunked” across multiple Batch processes.
 - Assume a batch job referencing 8000 reports is executed with 4 batch report processes is configured. The work would be chunked across 8 Batch Report Service threads.





Report Server: 32-bit vs. 64-bit

- Two Query Engines in Cognos Analytics: Compatible and Dynamic.
 - Compatible: 32-bit Report Server
 - Dynamic : 32-bit Report Server & 64-bit Report Server
- User end performance is comparable between the two Report Server Configurations.
- 64-bit Report Server could reduce the memory footprint on a system by using less report server processes.
- 32-bit Report Server process is Large Address aware. This means that the processes can grow to a larger size (up to 4 GB OS dependent).

The screenshot shows the IBM Cognos Configuration web interface. The title bar reads "IBM Cognos Configuration - dal-phastos.ottawa.ibm.com". The interface includes a menu bar (File, Edit, View, Actions, Help) and a toolbar with icons for save, refresh, and help. On the left is an "Explorer" pane showing a tree view of configuration categories: Local Configuration, Environment (selected), Configuration Group, Logging, File, IBM Cognos..., HeliosR7, Portal Services, Security, Authentication, Cognos, LDAP, Cryptography, and Cognos. The main area displays "Environment - Group Properties" as a table with two columns: Name and Value.

Name	Value
Dispatcher Settings	
* External dispatcher URI	http://dal-phastos:9300/p2pd/servlet/disp...
* Internal dispatcher URI	http://dal-phastos:9300/p2pd/servlet/disp...
Dispatcher password	*****
External JMX port	0
External JMX credential	*****
* Report Server execution mode	64-bit
Other URI Settings	
* Dispatcher URI for external applica...	32-bit
* Content Manager URIs	64-bit
* Content Manager URIs	http://dal-panther:9300/p2pd/servlet
Font Settings	
* Physical fonts locations	../bin/fonts
Physical fonts map	<click the edit button>
Fonts to embed (batch report servi...	<click the edit button>


Report Server: 64-bit



The memory footprint of Report Server can be reduced by varying the affinity, the session cache settings and the number of 64 bit Report Server processes. For example, these configurations handle the same number of work:

- 125 processes, 8 low and 2 high threads, session cache of 20
 - Memory: ~350MB per process = over 40GB
 - 25 processes, 40 low and 10 high threads, session cache of 100
 - Memory: ~700MB per process = under 20 GB
 - 10 processes, 100 low and 25 high threads, session cache of 250
 - Memory: ~1.3GB per process = under 15GB
-
- Calculate the session cache by taking the low and high affinity threads and multiple by two.
 - Affinity settings configured in the Cognos Administration Console.
 - Session cache is configured in the rsvpproperties.xml on disk in the configuration folder.

Gateway / Load Balancing



- Gateway, load balancing and security settings can have a huge impact on performance. Monitoring must be done to ensure a healthy system.
- Throughout the Cognos Analytics releases configuring a gateway has become easier with improved templates and documentation.
- Documentation:
 - https://www.ibm.com/support/knowledgecenter/SSEP7J_11.1.0/com.ibm.swg.ba.cognos.inst_cr_winux.doc/c_inst_all_config_gateway_intro.html
 - https://www.ibm.com/support/knowledgecenter/en/SSEP7J_11.1.0/com.ibm.swg.ba.cognos.inst_cr_winux.doc/c_config_web_server_intro.html
- Template files are located in the <cognos root>/cgi-bin/templates folder:

```
cognos_apache22_loadbalance.conf
cognos_apache24_loadbalance.conf
cognos_IHS85_loadbalance.conf
cognos_IHS9.conf
```

```
cognos_apache22_loadbalance_SSO.conf
cognos_apache24_loadbalance_SSO.conf
cognos_IHS85_loadbalance_SSO.conf
cognos_IHS9_SSO.conf
```



Cognos Analytics Port Usage

Default port settings for Cognos Analytics components

The following table lists the default ports and URI settings for IBM Cognos Analytics.

Table 1. Default port settings for Cognos Analytics components

Setting	Default Value	Description
Content Manager URI	<code>http://localhost:9300/p2pd/servlet</code>	The URI to Content Manager.
Gateway URI	<code>http://computer_name:port/bi/v1/disp</code>	The URI to the gateway.
Dispatcher URI (Internal, External)	<code>http://localhost:9300/p2pd/servlet/dispatch</code>	The URI to the dispatcher.
Dispatcher URI for external applications	<code>http://localhost:9300/bi/v1/disp</code>	The URI to the dispatcher.
Log server port	9362	The port used by the local log server.
Member synchronization port	4300	The local port used for network communication that transfers and synchronizes configuration information from one server to another.
Member coordination port	5701	The local port used for network communication for group coordination. This port is used to discover and join a group, and to maintain an up to date list of configuration group members.
Dataset Service port	9301	The local port that is used for inter-process communication. This port is assigned when Cognos Analytics is started for the first time. The port number is based on the Cognos Analytics dispatcher port plus 1. For example, 9300 +1 = 9301.

https://www.ibm.com/support/knowledgecenter/SSEP7J_11.1.0/com.ibm.swg.ba.cognos.ca_ig.doc/c_reviewthedefaultsettings.html#ReviewtheDefaultSettings

- Port allocation for the CognosServer and Dataset-Service is defined in the Cognos Configuration. By default the CognosServer port is 9300 and the Dataset-Service port is 9301.
- Other ports used by Cognos Analytics are outlined in documentation.

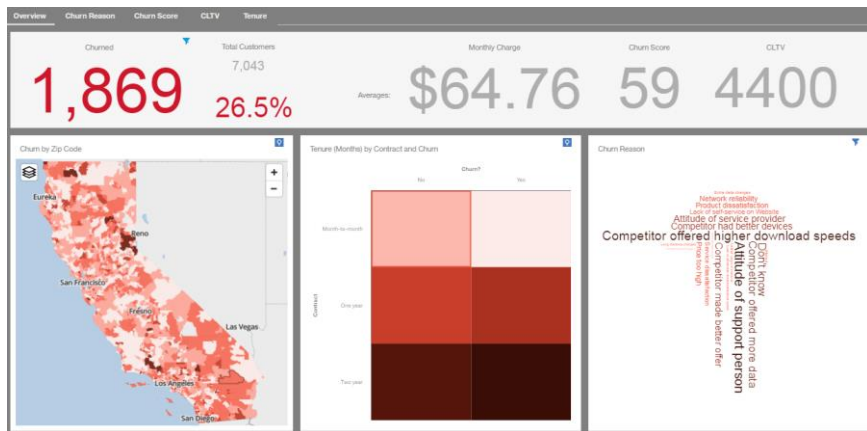
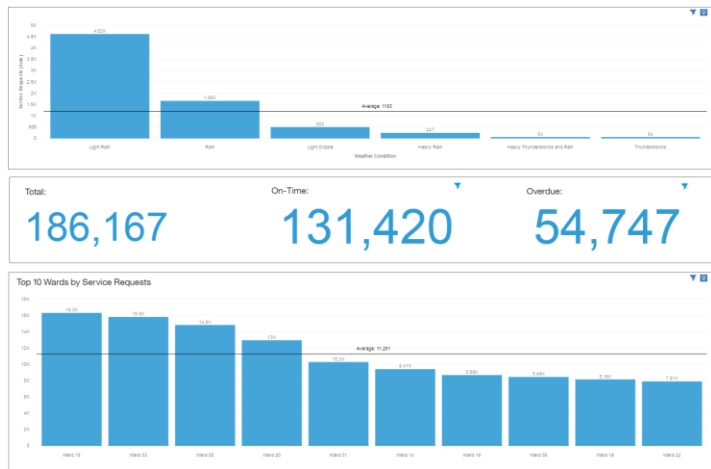
Cognos Analytics Performance

End User View

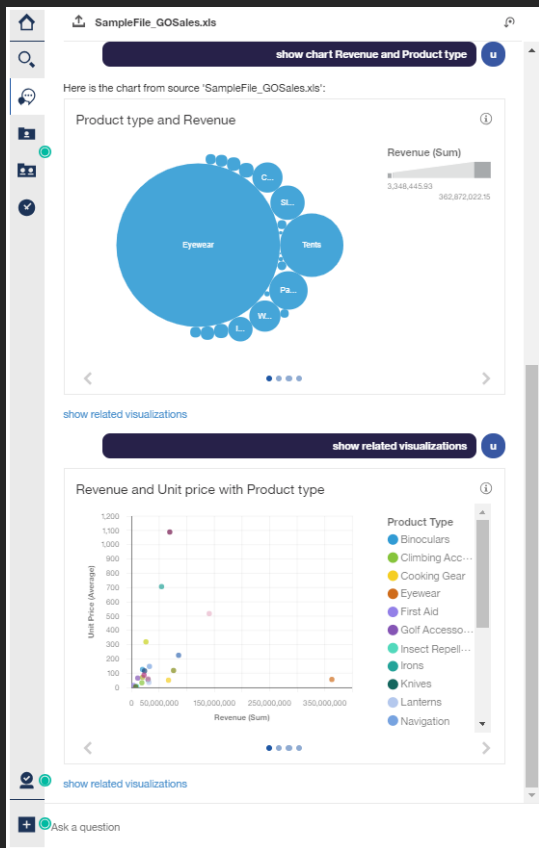


Dashboards and Stories

- Dashboards and Stories are tested against content that is created off of datasets, modules, packages and uploaded files.
- Dashboards and Stories use only the java processes. They do not use the BIBusTKServerMain process.
- Dashboards and Stories include smarts, predictive, conversation capabilities and annotations.
- Dashboards can be saved as Stories.
- Dashboard and Story specifications changed between CA 11.0 and CA 11.1. To ensure optimal performance resave all dashboards and stories when upgrading.



Conversation / Assistant

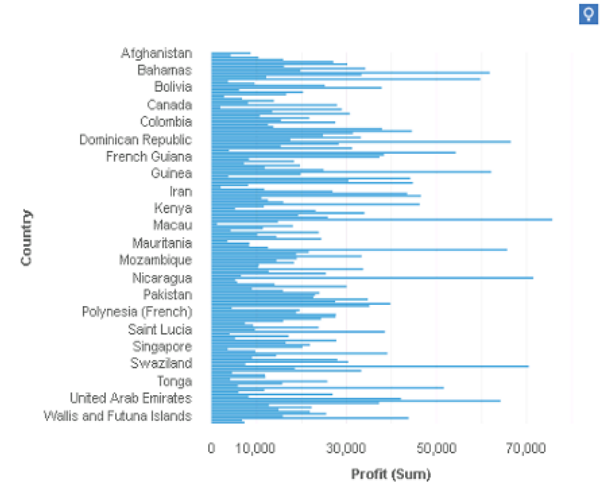
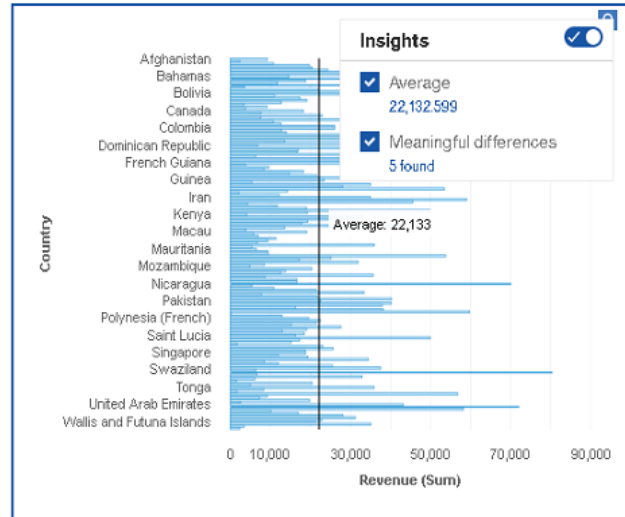


- The Conversation Assistant can help you drill into your data without leaving the workspace.
- The Assistant is available in dashboarding, storytelling and explorations.
- Content produced by the Assistant can easily be added to the workspace by dragging the information to the desired location.

Smart Annotations / Insights



- Smart Annotations can be enabled at the widget level for a dashboard, story or exploration widget.
- Enable the feature by activating the Insights option for the widget.
- The insights available are data and chart specific.
- By enabling the insights additional calls are made which can affect the performance of each widget. The expected degradation varies on the content, insights being queried and load already on the CA environment.






Exploration

- Explorations allow you to see smart relationships in data and build content useable by Dashboards and Stories.
- Explorations allow you to compare information side by side easily.
- Explorations use only the java processes. They do not use the BIBusTKServerMain process.

The screenshot displays a data exploration tool interface. On the left, a 'Cards' sidebar lists five visualization options: 1. A heatmap titled 'Tenure (Mont...ct and Churn', 2. A dot plot titled 'Tenure', 3. A line chart titled 'Total Charges...t and Tenure' with an average of 32.37, 4. A bar chart titled 'Tenure', and 5. A bar chart titled 'Tenure by Contract'. The main workspace is titled 'Explore relationships in your data' and includes a search bar for 'Field of interest' (set to 'Churn Value'), a filter for 'unique fields' (set to 2), and a list of fields: 'Churn Value', 'Monthly Charges', 'Total Charges', 'Churn Score', and 'CLTV'. Below this is a network graph with 'Churn Value' as the central blue node (Field of interest) and other green nodes (Fields) including 'Contract', 'Churn Reason', 'Churn Label', 'Churn Score', 'Payment Method', 'Tech Support', 'Online Security', 'Online Backup', and 'Internet Service'. A legend on the right identifies blue as 'Field of interest' and green as 'Field'. A 'Create a visualization' panel in the center offers instructions to 'Drag and drop fields or choose visualization type' and a 'Choose a type' button.

Compute Service / Flint



- The Compute Service JVM process is a sub-process launched and managed by the Query Service (dataset-service) JVM.
- The Compute Service process gets called when Parquet related queries are issued in order to improve performance and caching. If the same query is called multiple times and already exist, then a cached version will be returned instead of executing the query again.


Reporting

	18062512_baseline (309000)			1811151250_cur (317480)		
	Σ Average (s)	%Diff	90percentile	Σ Average (s)	%Diff	90percentile
1user_ResultDB_Off1MillionrowDataset_html_dqm	108.91		108.91	81.70	24.98% (27.21s)	81.70
1user_ResultDB_Off2MillionrowDataset_html_dqm	177.03		177.03	131.68	25.62% (45.36s)	131.68
1user_ResultDB_Off4MillionrowDataset_html_dqm	304.14		304.14	243.95	19.79% (60.20s)	243.95

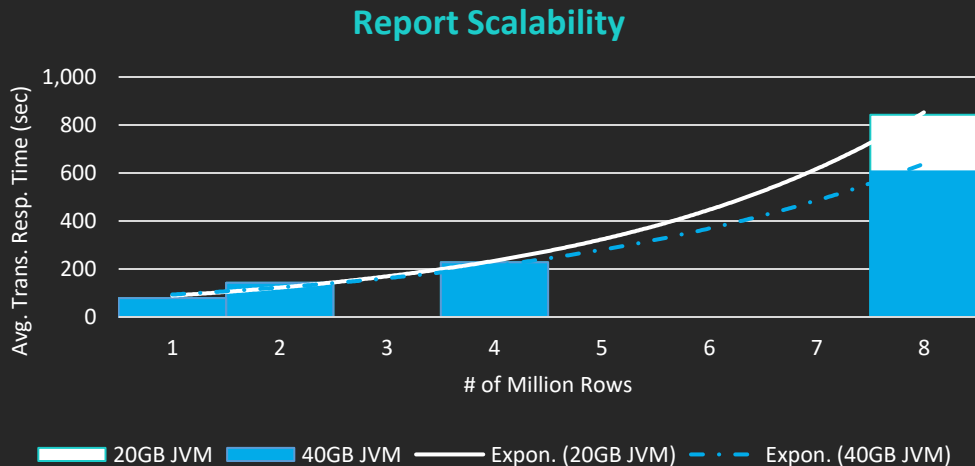
Dashboards

	18062512_baseline (309000)			1811151250_cur (317480)		
	Σ Average (s)	%Diff	90percentile	Σ Average (s)	%Diff	90percentile
1user_dashboardDataSet1Millrows_resultdb_same3Widget_har	21.86		21.86	4.83	77.92% (17.03s)	4.83
1user_dashboardDataSet2Millrows_resultdb_same3Widget_har	35.31		35.31	4.39	87.56% (30.91s)	4.39
1user_dashboardDataSet4Millrows_resultdb_same3Widget_har	64.03		64.03	4.30	93.29% (59.73s)	4.30
1user_dashboardModule1Millrows_resultdb_same3Widget_har	21.25		21.25	6.11	71.25% (15.14s)	6.11
1user_dashboardModule2Millrows_resultdb_same3Widget_har	35.31		35.31	4.84	86.31% (30.47s)	4.84
1user_dashboardModule4Millrows_resultdb_same3Widget_har	58.40		58.40	4.66	92.02% (53.74s)	4.66


Compute Service / Flint (Reporting)



- The number of rows in a dataset being brought back by a report can influence the overall scalability.
- If the scalability isn't predictable then the Query Service JVM may need to be increased to accommodate the data being retrieved.
- Below a loss of scalability is seen between 4 million and 8 million rows for a 20GB configuration.
- By increasing the Query Service JVM scalability is improved.




Compute Service / Flint (Dashboards)




- Scalability with dashboards in CA 11.1 is influenced by the user load more than the amount of rows in a dataset. This is due to the number of rows being clipped at 10000 for each widget on the dashboard.
- As user load increases response times increase predictably. If times do not increase predictably this could be due to the JVM being exhausted.

3 Widget DataSet Dashboard Scalability Warm Cache Results





Dataset vs. Database

Report and dashboard without a dataset



Report and dashboard with a dataset




- Building on top of the existing Cognos reporting platform, the dataset feature produces a data snapshot from a package, an external file or even a module. The dataset provides an on-the-fly, self-serving capability without the need for a complex tool or technical expertise.
- A developer can author much more performant content against a dataset using a module, due to the high-efficient data storage, simplified query process and improved data cache in the Cognos dataset service.
- In the top figure, the first scenario delivers dashboard/report via dynamic generated queries directly against the data source, providing near “real time” data reporting. The second scenario results in much faster performance via dataset with a smaller data volume. Note the data freshness depends on how frequent the dataset is built.
- The dataset feature is perfectly designed situations where data does change but changes in a manageable / predictable manner.
- <https://www.ibm.com/communities/analytics/cognos-analytics-blog/creating-high-performance-dashboard-and-report-using-cognos-dataset/>

JVM Tuning Comparison



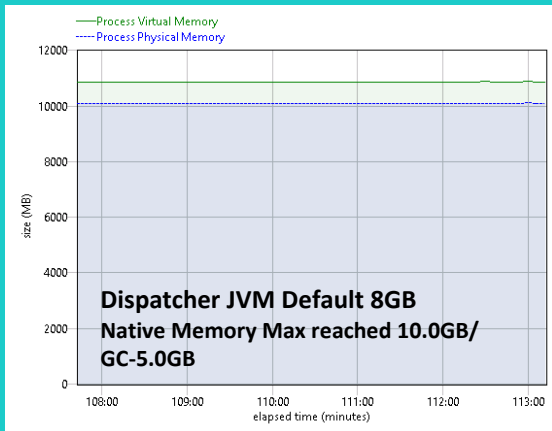
Cognos Analytics 11.1

JVM Tuning for the Cognos Server




The default JVM size for the cognosserver is 8GB in CA11.1. Is that good enough for my environment?

- After running 12 performance test cases ranging from 10 to 160 concurrent users executing file upload, dashboards or reporting the follow details that a 16GB configuration for the cognosserver process is ideal for this environment.
- The charts below show that the cognosserver JVM hit the native max setting of 8GB which could cause performance impacts due to frequent garbage collections and long GC pauses of 12.7% indicating the 8GB JVM heap setting is not adequate to handle the performance load. Ideally GC pauses should be less than 2%.






JVM Tuning for the Dataset-Service




The default JVM size for the dataset-service is 8GB in CA11.1. Is that good enough for my environment?

- After running 12 performance test cases ranging from 10 to 160 concurrent users executing file upload, dashboards or reporting the follow details that the 16GB configuration is ideal for this environment.
- The charts below show that the dataset-service JVM hit the native max setting of 8GB which could cause performance impacts due to frequent garbage collections and long GC pauses of 16.6% indicating the 8GB JVM heap setting is not adequate to handle the performance load. Ideally GB pauses should be less than 2%.




JVM Tuning for the Compute Service




The default JVM size for the compute service is 8GB in CA11.1. Is that good enough for my environment?

- After running 12 performance test cases ranging from 10 to 160 concurrent users executing file upload, dashboards or reporting the follow details that the 8GB max configuration is ideal for this environment.
- The Compute Service / Flint JVM of 8GB shows to be optimal for these performance scenarios with the default 8GB configuration with GC pauses of 0.79%.



JVM Tuning Performance



User end performance may or may not be impacted by the changing size of the JVM.

- For heavy actions such as uploading files performance is more likely to be impacted by a JVM that isn't properly tuned.
- If the JVM has long GC pauses then all operations will be impacted but potentially only for the amount of time the GC is occurring. Having a GC collection rate of less than 2% is ideal for all JVMs.

System Performance

Conformance



- Ensuring you are meeting the minimum conformance for Cognos Analytics is key to having an optimal system.
- Conformance points can be found at:

<http://www-01.ibm.com/support/docview.wss?uid=swg27047186>

11.1.x	11.0.0	10.2.2	10.2.1	10.2	10.1.1	10.1
--------	--------	--------	--------	------	--------	------

Please select an IBM Cognos Analytics 11.1 offering:

Select one

Cognos Analytics on Premises 11.1.x

Go to the table for your product release and then click a link to view a report.

Cognos Analytics on Premises 11.1.7 (LTS*)

* **Note:** Version 11.1.7 of IBM Cognos Analytics is a Long Term Support (LTS) release.

Requirements by type	Requirements by platform	Supplementary information
<ul style="list-style-type: none">• Operating Systems• Software (including application servers, data sources, and web browsers)• Hardware• Hypervisors	<ul style="list-style-type: none">• AIX• Linux• Mobile OS• Windows	<ul style="list-style-type: none">• Supported and tested client drivers 11.1.7 FP2 [Relational] [OLAP]• Supported and tested client drivers 11.1.7.0 [Relational] [OLAP]

Operating System Tuning: ulimits



- It is important for Cognos Administrator and System Administrator to be aware of any impact on setting the proper resource (particularly ulimit values) on Linux operating systems as it can affect how Cognos Analytics performs. Often the symptoms are inconsistent and vary depending on the workload being executed.
- The ulimit command on Unix/Linux sets or reports user process resource limits.
- On Unix and Linux operating systems, ulimit is configurable and usually set too low for Cognos Analytics by default.
- These limits are categorized as either soft or hard. With the ulimit command, you can change soft limits, up to the maximum set by the hard limit (set by root). It is either displayed as 'open files' or 'nofiles'.
- <http://www-01.ibm.com/support/docview.wss?uid=swg21067352>

```
ulimit -n      specifies the limit on the number of file descriptors a process may have
ulimit -a      displays the current resource limits
ulimit -Ha     displays the hard limit for a given resource (configurable by the root user)
To set, use the command ulimit -n <value>
```

Test #	Soft limit (ulimit -Sa)	Hard limit (ulimit -Ha)	Security	Result
1	1024	8192	Anonymous	With low values setting, some process won't start (e.g BmtProviderMain).
2	1024	8192	LDAP	Additional errors on Content Manager. Key functionality such as importing deployment fails.
3	8192	8192	LDAP	Started OK, quick user load test can be performed such as dashboard and reports.

Operating System Monitoring

- Important to use tools that allow for unattended monitoring of resource utilization over time. For example:
 - PerfMon for Windows
 - Nmon on AIX/Linux (use Nmon Analyser to process Nmon output)
- Read the Cognos Analytics documentation for any OS specific settings that may need to be applied.
- Bottlenecks in the system resources can lead to frustration in the Analytics community due to:
 - Inconsistent performance
 - Unexpected error messages
- As hardware and software evolve, system bottlenecks tend to shift. The four most common system bottlenecks are:
 - CPU
 - Memory for both the entire system and key Analytics processes
 - Network utilization
 - Disk for read, writes, and waits



Operating System Monitoring

If CPU is a bottleneck:

- Shift Analytics services to other servers in the system that may have available CPU.
- Consider adding additional CPU resources or an additional server.
- Monitor Run Queue.

If memory is a bottleneck:

- Turn off services and processes to free up memory.
- Add more memory to the system or shift Analytics services to servers with more available RAM.
- Do not rely on Virtual Memory. Absolute performance killer!

If network appears to be a bottleneck:


- Check that the NIC card is using the full bandwidth available.
- Ensure the server resolves localhost locally first and not to the DNS first (netsh.exe).
(netsh.exe advbase set "Local Settings" "LocalHostResolution" "Local")
- Ensure routers in the Analytics system are not overtaxed.

If disk might be a bottleneck:

- Check that file system logging is turned off or minimized (mount).
- Consider using fast storage to host disk intensive Analytics services.



Monitoring Cognos Analytics



Health Check Service

<dispatcher URL>/bi/health


```
▶ account-service: (-)
▶ admin-service: (-)
▼ bi:
  healthy: true
  message: "All is well."
  ▼ sysInfo:
    uptime: "04:06:08.866"
    totalMemory: "4.6 GB"
    usableDiskSpace: "215.0 GB"
  ▼ gcStats:
    collectionTime: "00:00:01.891"
    ▼ memoryPoolNames:
      0: "Java heap"
      name: "Marksweepcompact"
      collectionCount: "1"
    processors: 4
    freeMemory: "1.0 GB"
    maxMemory: "8.6 GB"
    service: "/bi"
▶ bi-api: (-)
▶ collaboration-service: (-)
▶ config-service: (-)
▶ content-service: (-)
▼ dataset-service:
  healthy: true
  ▼ message:
    service: "dataset-service"
  ▼ metrics:
    NumberOfReceivedRequests: 0
    NumberOfCriticalFailures: 0
    NumberOfSuccessfulRequests: 0
  status: "running"
▶ flipper-service: (-)
▶ geo: (-)
▶ glug-service: (-)
▼ modelingservice:
  healthy: true
  message: "File uploading is good!"
▶ notification-service: (-)
▶ predict-suggestion: (-)
▶ report-service: (-)
▶ scheduling-service: (-)
▼ search:
  healthy: true
  message: "Search-Service is running ok!"
▶ smarts: (-)
▶ smarts-conversation/api: (-)
▶ smarts-explore/api: (-)
▶ smarts-insights/api: (-)
▶ smarts-modeling/api: (-)
▶ smarts-visualization-recommender/api: (-)
▶ ui/v1: (-)
▶ user-profile-service: (-)
▶ wlpApps: (-)
```

- The health of a dispatcher can be monitored by hitting the **<dispatcher>/bi/health URL**.

- The status of services such as the bi, dataset-service, modelingservice and search can be tracked by monitoring the URL.

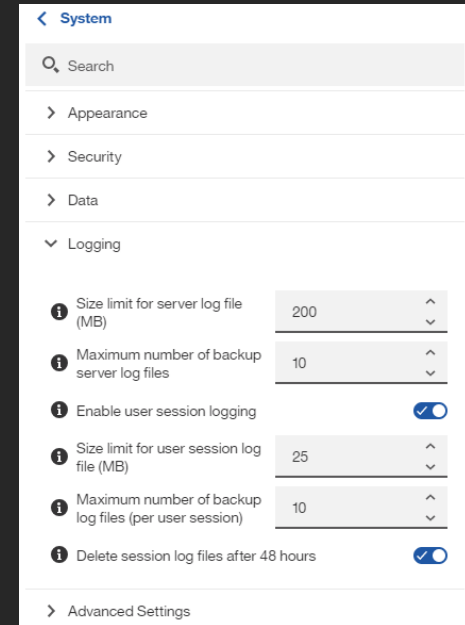
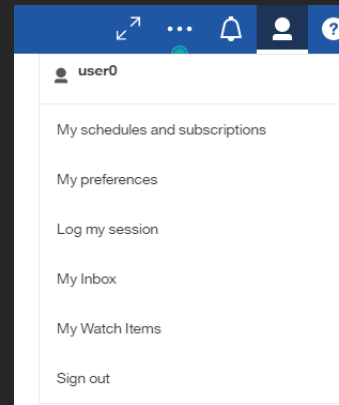
- For an environment with multiple dispatchers it is best to hit each dispatcher URL individually to ensure the status is accurate.

Cognos Analytics Session Logging



Cognos Analytics logging changed in CA 11.0.7.

- The purpose of this logging change is to help in diagnosing complex problems by collecting related logs from easily accessible web-based menus.
- Before enabling diagnostic logging, the Cognos Administrator can review and update size limits and the maximum number of files to keep.
- Enable system logging through the CA portal at Manage > Configuration, select the “Diagnostic Logging” tab.
- User session logging can be enabled by the user in their preferences.



Cognos Analytics Basic Logging

- Cognos Analytics by default is set to a service level logging of Minimal.
- Minimal, Basic, Request, Trace and Full are the options available for service level logging. The logging level is configured in Cognos Administration.
- By enabling Basic logging for targeted services details on the server side throughput can be monitored and tracked.
- Service level logging is written to the < Cognos root >/cogaudit.log file.

cogaudit.log basic logging example:

```
192.168.40.1:9300 40304 2018-12-13 11:06:08.492 -  
5 1346...BBC8 68a7b7b5:167a84b3be9:4f1c 2614 DSS 0 3  
Audit.RTUsage.qs.DSS Execute Exploration  
QueryService Success <parameters></parameters>
```

<input type="checkbox"/>	Category	Name	Value	Default
<input type="checkbox"/>	Logging	Audit logging level for annotation service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for agent service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for batch report service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit the native query for batch report service	<input type="checkbox"/>	Yes
<input type="checkbox"/>	Logging	Audit logging level for the Content Manager Cache Service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for Content Manager service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for the dispatcher	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for delivery service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for event management service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for graphics service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for human task service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for interactive discovery visualization service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for job service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for mobile service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for metadata service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for monitor service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for presentation service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for query service	Minimal	Yes
<input type="checkbox"/>	Logging	Enable query execution trace	Minimal	Yes
<input type="checkbox"/>	Logging	Enable query planning trace	Basic	Yes
<input type="checkbox"/>	Logging	Generate comments in native SQL	Request	Yes
<input type="checkbox"/>	Logging	Write model to file	Trace	Yes
<input type="checkbox"/>	Logging	Audit logging level for report data service	Full	Yes
<input type="checkbox"/>	Logging	Audit logging level for relational metadata service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for report service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit the native query for report service	<input type="checkbox"/>	Yes
<input type="checkbox"/>	Logging	Audit logging level for repository service	Minimal	Yes
<input type="checkbox"/>	Logging	Audit logging level for system service	Minimal	Yes



Interactive Performance Assistant (IPA)

- Enable and use Interactive Performance Assistant to view query and rendering performance on report objects in Cognos Analytics.

- <https://www.ibm.com/communities/analytics/cognos-analytics-blog/interactively-view-report-performance/>

Run options [Close]

Paper size: (Default)

Paper orientation: (Default)

Data mode: All data

Language: (Default)
Afrikaans
Afrikaans (South Africa)
Albanian
Albanian (Albania)
Amharic

Rows per page: []

Prompt

Include accessibility features

Enable bidirectional support

Include performance details

[OK] [Cancel]

Execution Time: 664 ms

California Cites: 292 Customers: 7,128

Execution Time: 193 ms

PROFIT_MARGIN_AMT_B	1999												Total(YEAR_L)
	01	02	03	04	05	06	07	08	09	10	11	12	
Pool & Spa Group	4,069,895.06	2,131,656.67	4,623,823.38	3,861,686.84	5,751,834.55	11,382,447.56	5,986,683.89	2,978,537.87	2,038,438.49	4,512,412.08	4,900,284.07	3,660,655.35	55,898,355.81
Surface Water Group	677,368.67	850,739.12	818,282.41	1,096,035.05	1,634,824.18	2,728,781.11	914,606.88	1,122,365.86	1,735,449.65	1,464,866.26	1,538,835.98	1,883,776.18	16,465,931.35
Total(CUSTOMER_CAT_L)	4,747,263.73	2,982,395.79	5,442,105.79	4,957,721.89	7,386,658.73	14,111,228.67	6,901,290.77	4,100,903.73	3,773,888.14	5,977,278.34	6,439,120.05	5,544,431.53	72,364,287.16

Execution Time: 29638 ms

Execution Time: 30699 ms

Aug 25, 2017

Total Execution Time [Request Id: yjvhGsywM44H8qMjs8lyq4j4lCq22y4lqjG]: 47924 ms

Monitoring Tools: GC Logs



JVM Memory and GC Policies:

- Enabling GC logging is a low impact method of measuring JVM sizing and Garbage Collection policies. Undersized JVMs can lead to OOM situations or excessive garbage collections and high JVM pause times.
 - For Dispatcher and CM, edit <cnos root>/bin64\bootstrap_wlp_<OS>.xml and add the following line to the Java arg list:

```
<param condName="{java_vendor}" condValue="IBM">-Xverbosegclog:../logs/disp_gc.log</param>
```

- For CGS on Windows edit <cnos root>\webapps\p2pd\WEB-INF\services\cgsService.xml on Windows and add the following line to the JVM arguments (in two locations in the file):

```
<child-proc-cmd>Xverbosegclog:../logs/cgs_gc.log</child-proc-cmd>
```

- For CGS on Unix edit <cnos root>\cgsServer.sh on UNIX and add -Xverbosegclog to the \$JAVA_OPTS JVM argument line.
- For the Query Service, GC logging is on by default in a file named dq_verbosegc_%timeStamp%.log
- For the Compute Service, GC logging is on by default in a file named flint_verbosegc_%timeStamp%.log



Load Balancing Information

<gateway URL>/ibmcognos/balancer-manager

- When using the Cognos Analytics template gateway configuration files a “balancer-manager” is provided by default.

```
# UI to monitor/configure load balancer
<Location /ibmcognos/balancer-manager>
  SetHandler balancer-manager
</Location>
```

- By accessing the <gateway URL>/ibmcognos/balancer-manager URL a gateway level health check of the CA environment is seen.

- Depending on if you are using Apache or HTTP Server the balancer-manager information is slightly different.

Load Balancer Manager for localhost

Server Version: IBM_HTTP_Server/8.5.5.11 (Win32)
Server Built: Oct 31 2016 10:05:47

LoadBalancer Status for balancer://mycluster

StickySession	Timeout	FailoverAttempts	Method
-	0	5	bybusyness

Worker URL	Route	RouteRedir	Factor	Set	Status	Elected	To	From
https://drv-saxony.canlab.ibm.com:9300	1		1	0	Ok	49971250	185G	226G
https://vott-stab03.canlab.ibm.com:9300	2		1	0	Ok	39356644	147G	177G
https://duv-pair.canlab.ibm.com:9300	3		1	0	Ok	53296228	204G	246G
https://drv-pharate.canlab.ibm.com:9300	4		1	0	Ok	38861984	140G	172G
https://dny-partner.canlab.ibm.com:9300	5		1	0	Ok	40033510	149G	182G
https://dal-seabed.canlab.ibm.com:9300	6		1	0	Ok	47792225	177G	216G



Load Balancing Troubleshooting

- To configure Cognos Analytics so you can see which application server is processing the requests.
- Once enabled all CA HTTP responses will have a header called X-CA-Server with the dispatcher details.
- Enabling this feature may have a slight impact on performance

Under Manage > Configuration > System > Advanced Settings add the following key-value pair: CAServerHeader : true

```
HTTP/1.1 200 OK
Date: Tue, 13 Nov 2018 15:08:27 GMT
X-CA-Affinity: 23046953
Cache-Control: private,must-revalidate
Expires: Thu, 01 Jan 1970 00:00:00 GMT
Strict-Transport-Security: max-age=31536000
X-CA-Server: dal-phastos.ottawa.ibm.com:9300
Etag: "IBMWed Dec 31 19:00:00 EST 19690explore-
Content-Type: application/json
X-Compressed-By: BICompressionFilter
Content-Language: en-US
Content-Length: 53076
Keep-Alive: timeout=10, max=61
Connection: Keep-Alive
```

Monitoring Tools: IBM JVM Tools

www.ibm.com/developerworks/java/jdk/tools

Health Center

•Monitors the status of a running application to help you to:

- Optimize application performance
- Improve application stability and uptime
- Reduce system resource usage
- Reduce the time to resolve problems
- Drive down development and maintenance costs

Garbage Collection and Memory Visualizer

Analyzes output from garbage collection logs to:

- Visualize the memory usage and garbage collection activity of your application
- Provide recommendations for tuning your application

Dump Analyzer

Analyzes a formatted system dump to:

- Produce a report
- Provide suggestions on how to resolve the problem

Memory Analyzer

Analyzes system level dumps and heap dumps to:

- Help you understand what's in your Java heap
- Assist with debugging application issues

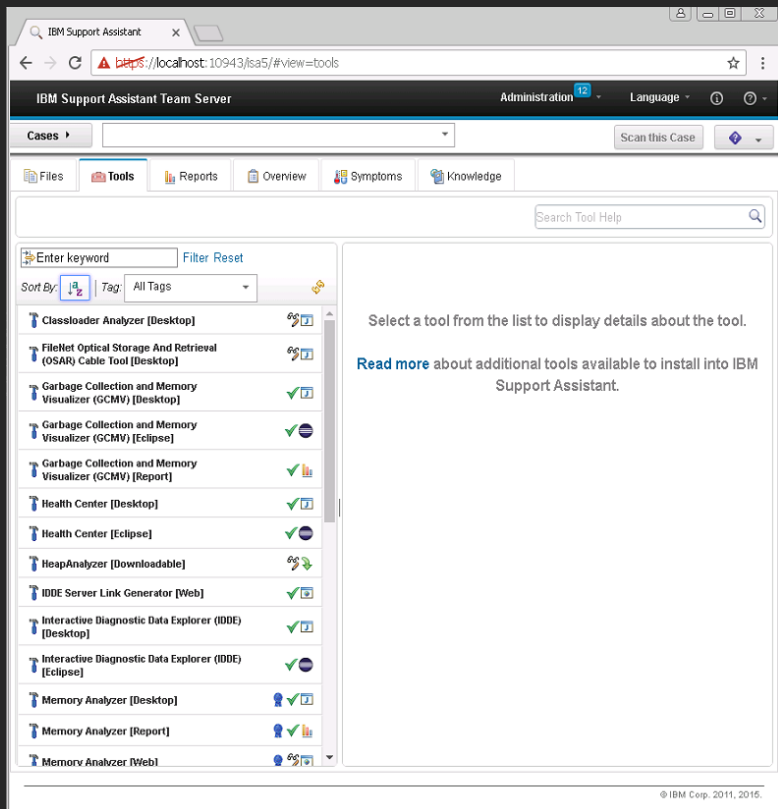
Interactive Diagnostic Data Explorer

Explores system core files and Java dumps to:

- Provide a visual representation of your core files
- Help you investigate problems interactively



Monitoring Tools: IBM Support Assistant



- The desired toolset are 'IBM Monitoring and Diagnostic Tools for Java'.

IBM Monitoring and Diagnostic Tools for Java™ - Dump Analyzer

IBM Monitoring and Diagnostic Tools for Java™ - Garbage Collection and Memory Visualizer for ISAv4

IBM Monitoring and Diagnostic Tools for Java™ - Health Center


IBM Monitoring and Diagnostic Tools for Java™ - Interactive Diagnostic Data Explorer

IBM Monitoring and Diagnostic Tools for Java™ - Memory Analyzer

IBM Monitoring and Diagnostic Tools for Java™ - Memory Analyzer 64bit


- Quick steps to use the tool:
 - Highlight the 'Garbage Collection and Memory Visualizer' tool and hit 'Launch' and browse to the garbage collection log using the 'Remote Artifact Browser'.
 - Arguably the most useful view is 'Heap size', 'Pause time', and 'Used heap (after collection)' chosen from the File menu 'VGC pause' and 'VGC heap'.

IBM Support Assistant: GC Logging Analysis




- Classic example of the IBM Support Assistant helping determine ‘ideal’ JVM size:
 - Left graph shows a properly tuned JVM with GC occurring approximately every 2 hours with a pause time of under 35 seconds.
 - Right graph illustrates a JVM running an undersized Java Heap. GC occurs every 10 minutes with a pause time of over 40 seconds.

Tuned JVM (using Xms/Xmx = 100 GB)



Small JVM (using Xms/Xmx = 50 GB)



Healthy Analytics System

Summary

Concurrent collection count	1
Forced collection count	0
GC Mode	genccon
Global collections - Mean garbage collection pause (ms)	543
Global collections - Mean interval between collections (ms)	1534407
Global collections - Number of collections	1
Global collections - Total amount tenured (MB)	421
Largest memory request (bytes)	6813808
Number of collections triggered by allocation failure	16
Nursery collections - Mean garbage collection pause (ms)	469
Nursery collections - Mean interval between collections (ms)	163418
Nursery collections - Number of collections	16
Nursery collections - Total amount flipped (MB)	1615
Nursery collections - Total amount tenured (MB)	433
Proportion of time spent in garbage collection pauses (%)	0.62
Proportion of time spent unpaused (%)	99.38
Rate of garbage collection (MB/minutes)	127