

Clustered Database Reporting Solution utilizing Tivoli

Clea Zolotow, Senior Technical Staff Member
Transition and Transformation Specialty Service Area
Global Technology Services
Internet zolo@us.ibm.com

Hariet von Moltke
Global Technology Services
Internet: hvonmolt@us.ibm.com

Joe Caldwell

IBM Global Services Enterprise Automation - Senior IT Architect
Strategic Outsourcing Delivery, Delivery Technology & Engineering
Internet: jcald@us.ibm.com

Tony Hunt, Distinguished Engineer
Transition and Transformation Specialty Service Area
Global Technology Services
Internet: tonyhunt@us.ibm.com

Table of Contents

Table of Contents	3
Table of Figures	4
Abstract	5
Scope of Database Monitoring.....	5
Tivoli Enterprise Portal/Netcool vs. Tivoli Database Warehouse	7
The Raw Database Problem.....	12
Database Clustering Options.....	13
DB2.....	13
Oracle RAC.....	15
Oracle RAC performance monitoring.....	17
Sybase	24
MS-SQL.....	26
Category 1: Performance reports	27
Category 2: Utilization reports.....	27
Category 3: Availability reports.....	27
Informix	28
Informix ITCAM Older Solution.....	30
GSMA Newer Informix Scripting Element	31
CONCLUSION.....	34
Appendix A: TPC Reports (non-Clustered DB)	35

Table of Figures

Figure 1: DB Clustering Monitoring Solution Overview.....	7
Figure 2: Tivoli Common Reporting for the Various Agents.....	8
Figure 3: Tivoli Monitoring Overview.....	9
Figure 4: Current Customer Monitoring Environment	11
Figure 5: Raw Database Example.....	12

Abstract

Each different midrange database type, such as DB2, Oracle RAC, Sybase, MS-SQL (i.e., SQL-Server), and Informix, has a different architecture leading to a different reporting methodology. This paper shows architecture for utilizing Tivoli's Composite Application Manager (ITCAM) to provide reporting for many clustered databases. It shows which databases and which architectures are supported and which require custom tooling. Standard reporting is then identified for each supported database system. At IBM, this is part of the Global Server Management Architecture (GSMA) and part of the Netcool Event Management System.

Scope of Database Monitoring

In this study, we are including the following databases:

- DB2
- Oracle RAC
- Sybase
- MS-SQL (i.e., SQL-Server)
- Informix

Some of the ways to do database monitoring are as follows:

- Agent (Tivoli, updated with GSMA)
- Agent (Other, sometimes GSMA)
- CLI interface (scripted)
- CLI interface (manual, i.e., logon and create a script)
- Homemade Health check Scripts (such as those from the DMSL Tools Portal at IBM)

Below, we present a dual database monitoring solution that depends mostly on Tivoli ITCAM agents, whenever possible.

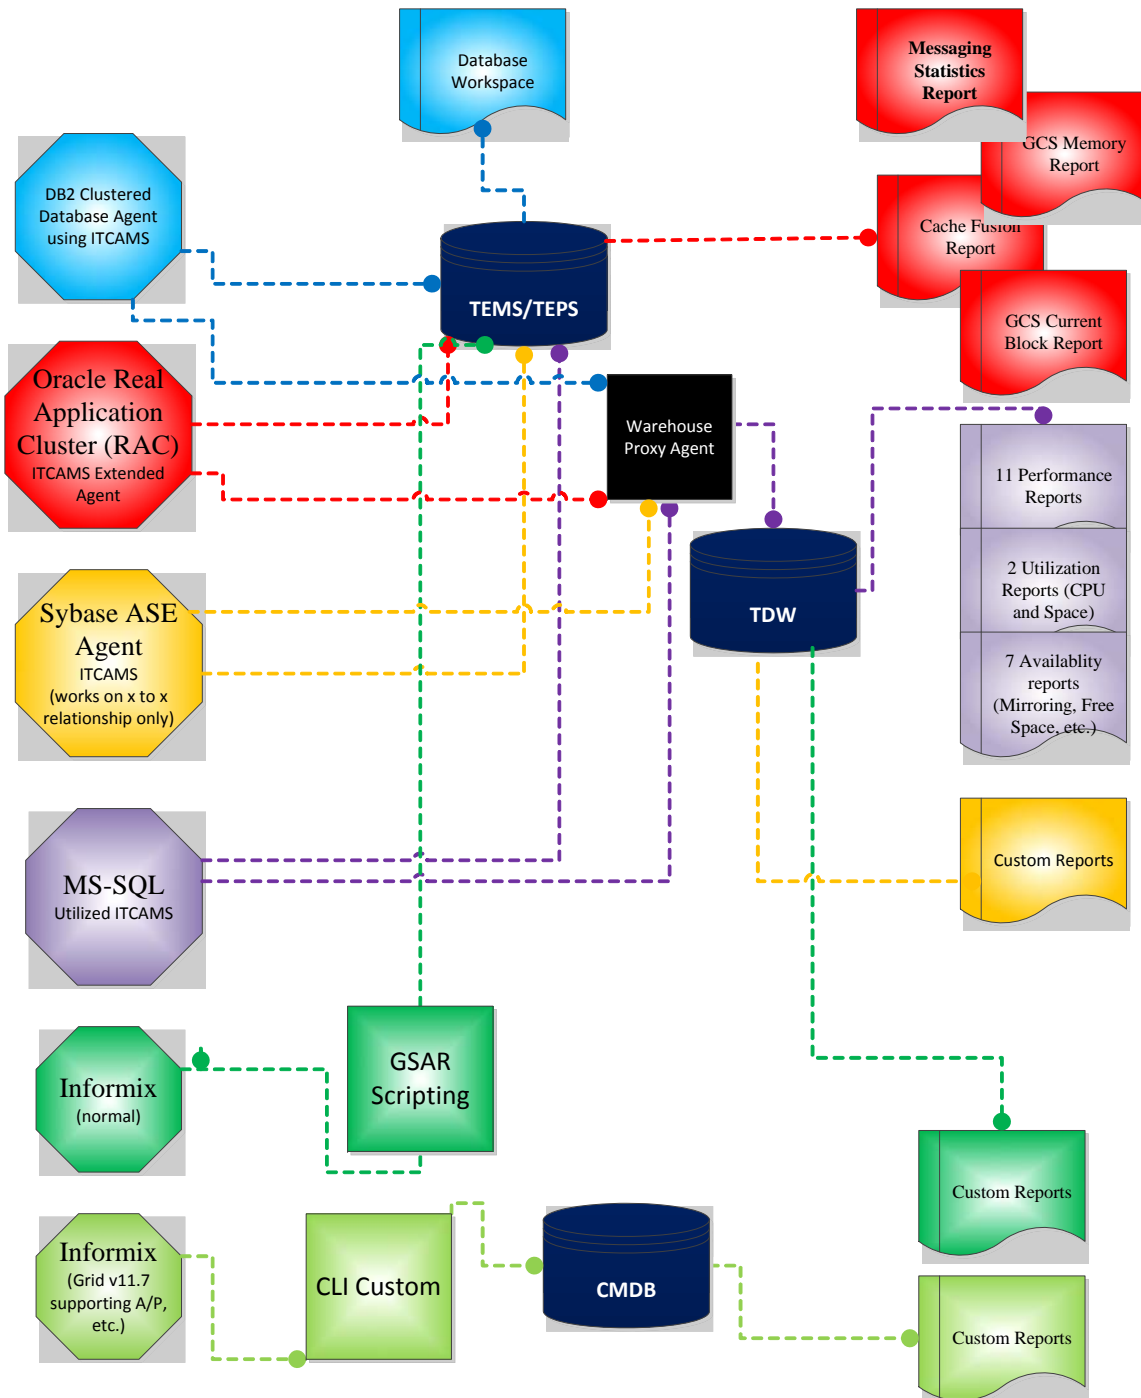


Figure 1: DB Clustering Monitoring Solution Overview

As many companies are already licensed for ITM as well as the ITCAM product line, we will look at these potential solutions first. CLI agents require that back end databases be created and the CLI be updated whenever systems change. When looking at the GSMA standards in IBM's Global Solutions Architecture Review (GSAR), many of these database types have monitoring profiles. Specifically: "GSMA provides IBM Tivoli Monitoring 6 (ITM6) monitoring solutions for Databases enabling the monitoring of resources on the following Relational Database Management System (RDBMS):

- DB2
- Oracle
- Microsoft SQL Server
- Sybase
- Informix

These solutions are based currently on the following products which are part of the ITM6/ITCAM family:

- ITCAM for Applications: which provides monitoring agents for DB2, Oracle (and Extended Oracle) and Sybase
- ITCAM for Microsoft Applications: which provides a monitoring agent for Microsoft SQL Server
- an OPAL solution which provides a Universal Agent for Informix monitoring (which has been superseded by the GSAR scripting)
- Custom agents developed by the GSMA BlueCARE team which provides Oracle monitoring.

Tivoli Enterprise Portal/Netcool vs. Tivoli Database Warehouse

Note that most database cluster reporting takes place live and is available in the Tivoli Enterprise Portal (TEP) or in the Netcool Web GUI. Alerting can be set up. Specific warehousing of historical data is made available to Tivoli Common Reporting (TCR) using the databases shown in the matrix below:

Database	Version
Oracle	Oracle V9.2, V10g Release 1, and V10g Release 2
DB2	DB2 V8.1 Fix Pack 10 and later, V8.2 Fix Pack 3 and later, V9.1 and fix packs
MSSQL	SQL Server 2000 Service Pack 4, SQL Server 2005 Service Pack 2, SQL Server 2008 Service Pack 1

Figure 2: Tivoli Common Reporting for the Various Agents

Other data elements can be loaded into existing Tivoli databases such as the Tivoli Data Warehouse (TDW) or the CMDB and reports can be created using Tivoli Common Reporting (TCR) which utilizes COGNOS. Today, there is a customizable linkage between the TEMS and the TDW, using the Warehouse Proxy Agent (on the RTEMS). Real-time event data loaded into the Netcool Web Gui will have to be utilized there or historical event data can loaded into the TDW using an adapter.

IBM Tivoli Monitoring environment
Architecture Overview

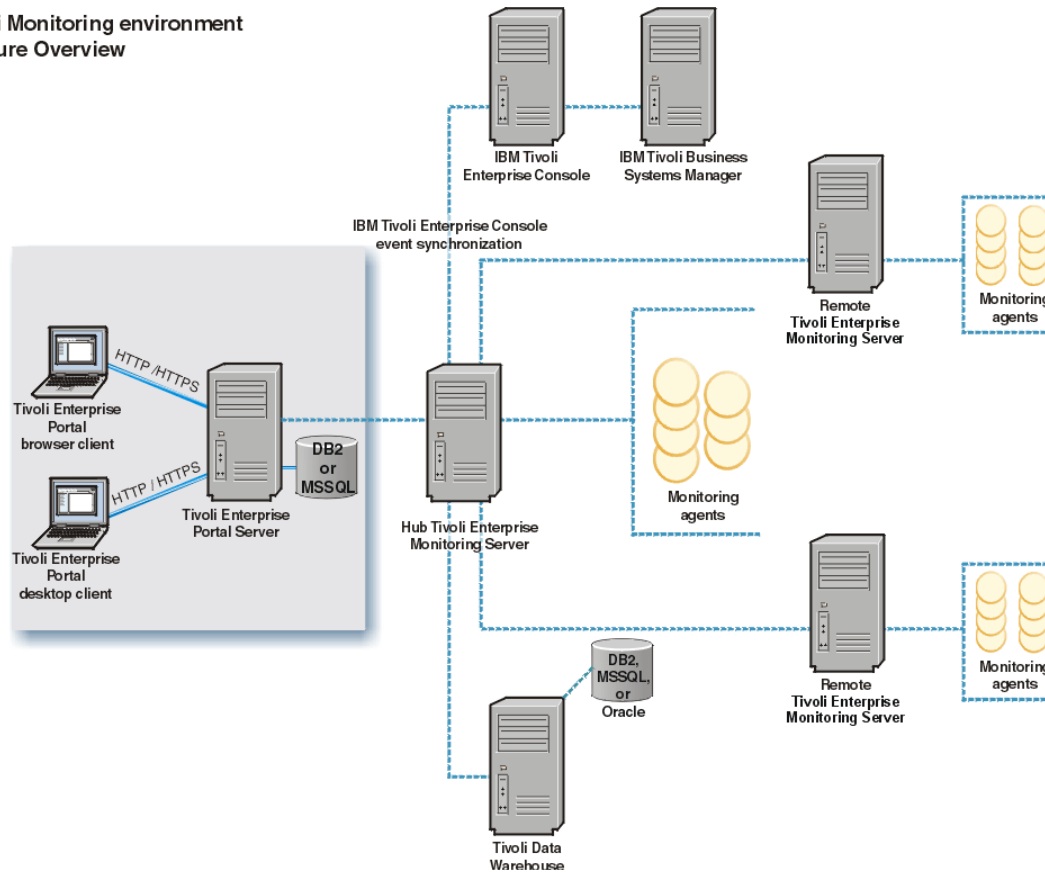


Figure 3: Tivoli Monitoring Overview

The Tivoli Monitoring environment consists of the following components:

- Hub TEMS – One of the Tivoli Enterprise Monitoring Servers (TEMS), which acts as a collection and control point for alerts received from the agents, and collects agents' capacity, performance and availability data. One server in each environment must be designated as the hub.

- Remote TEMs (RTEMS) - These monitoring servers are the collection and control points for performance, capacity and availability data and alerts. They are used for scalability purposes and report into the Hub TEMS
- Warehouse Proxy Agent (WPA) - can installed on the Hub TEMS or on each Remote TEMS in the environment. This agent is responsible for collecting and consolidating all historical data from the agents and storing the data into the Tivoli Data Warehouse.
- Tivoli Enterprise Portal (TEP) - The TEP provides visualization for Tivoli Monitoring. The TEP provides a graphical user interface (GUI) that includes navigation and visualization of databases, their status, key performance metrics and historical data.
- Tivoli Enterprise Portal Server (TEPS) - retrieves historical data for display in data views in the Tivoli Enterprise Portal. It retrieves short-term historical data from files stored on the monitoring agents. It retrieves long-term historical data from the Tivoli Data Warehouse.
- Tivoli Data Warehouse Server - stores data in the TDW database for historical analysis, trending, and reporting. Data collected by monitoring agents is stored on the agent at intervals in binary files as historical data. The historical data is sent from the monitoring agents to the Warehouse Proxy agent at a preset interval. The Warehouse Proxy agent sends the data it receives to the Tivoli Data Warehouse.
- The Summarization and Pruning agent retrieves detailed data from the warehouse database, aggregates or prunes the data, and returns the processed data to the warehouse. The Summarization and Pruning agent will create tables, indexes, and views as needed.

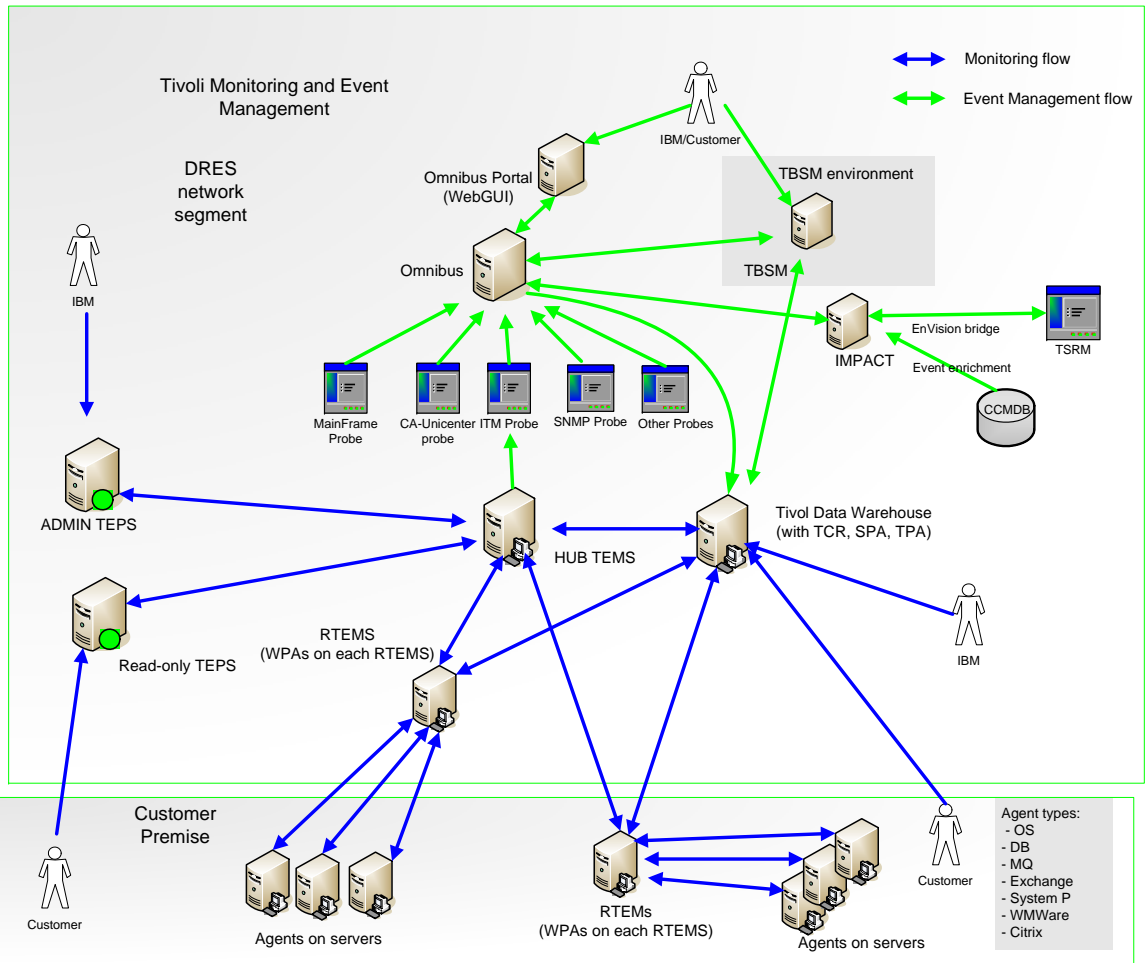


Figure 4: Current Customer Monitoring Environment

The Raw Database Problem

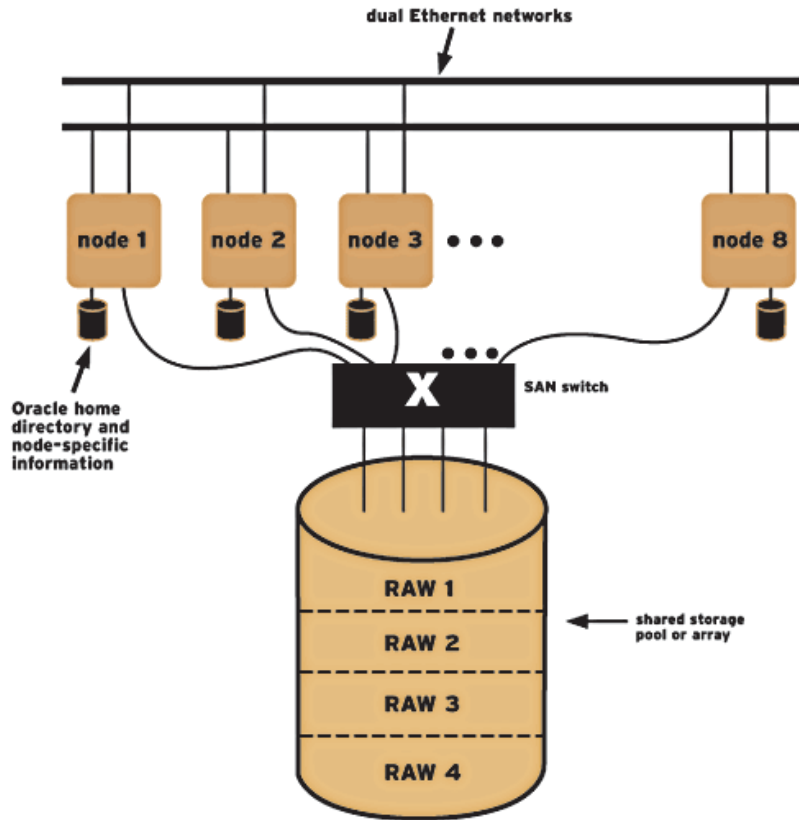


Figure 5: Raw Database Example¹

A raw database is defined as a raw device or raw disk that uses an unformatted disk partition. Raw devices do not have a file system mounted on them, and they are written to and read from character device drivers. Due to the absence of a regular file system such as JFS, UFS or NTFS, such partitions are considered to be raw or "uncooked". When a regular

¹ http://www.redhat.com/magazine/009jul05/features/gfs_oracle/#raw-devices

file system is used, the file system manages all I/O requests, each of which passes through the file system buffers. In other words, data is written to the file system buffers and then synchronized with the physical disks. When Oracle files such as data files, online redo log files, or control files are created on raw disk, Oracle takes responsibility for reading from and writing to that disk, since the file system is absent.

Unix uses device files to communicate with hardware components such as disk drives and tape drives. Each disk drive has corresponding entry in the /dev directory. A cooked disk may have both a block device file and a character device file associated with it. However, a raw disk has only a character device file. A character device file uses its own device drivers to read from and write to a disk and performs I/O operations in variable sizes. A block device file is dependent on file system buffers for performing I/O operations, and it reads from and writes to a disk in fixed-length size.”²

Since there is no file system, we have excluded raw databases from our clustering study.

Therefore, looking at specific database clustering that uses a file system, we start looking at the options for DB2, Oracle, Sybase, MS-SQL, and Informix.

Database Clustering Options

DB2

DB2 information can be found in Tivoli® Composite Application Manager (ITCAM) Agent for DB2³ documentation. Utilizing this documentation, we can set up agents for Microsoft Cluster Environments (i.e., Windows) and HACMP Environments (i.e., AIX). Specific, specialized setup instructions are included in the referenced documentation.

² <http://www.dbametrix.com/raw-device-oracle-tuning.html>

³ http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=%2Fcom.ibm.itcama.doc_6.2.4%2Fwelcome.htm

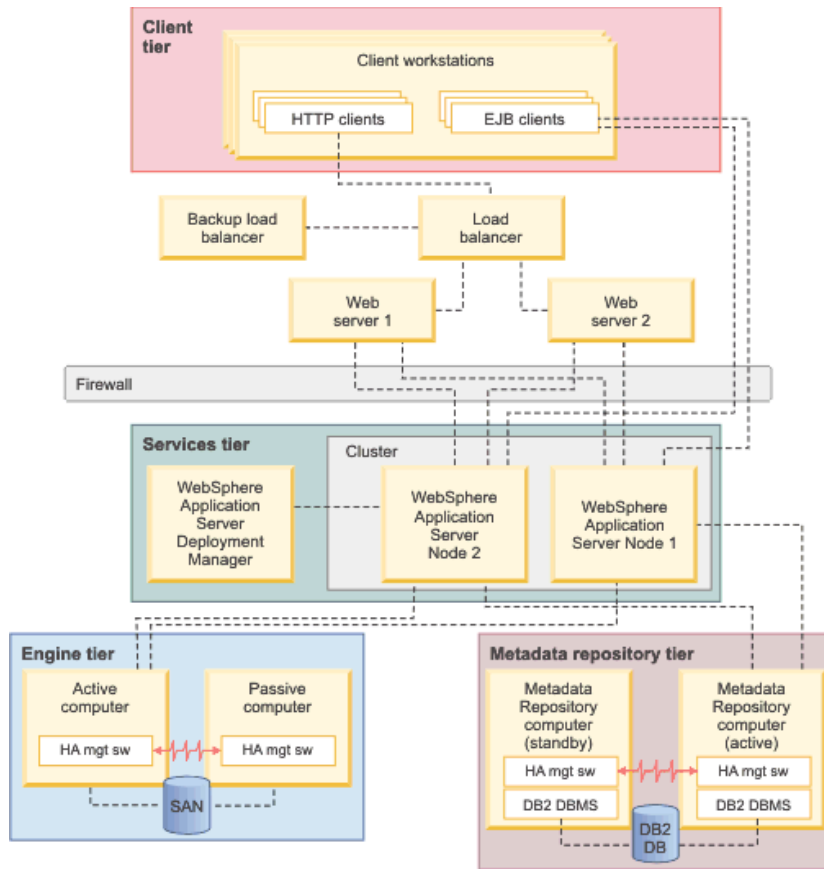


Figure 6: Clustered Metadata Repository Tier for DB2⁴

To provide high availability, clusters are set up in an active-passive configuration with a single active DB2 instance on one computer and one or more passive instances on the other computers, as shown in the figure above. If the DB2 instance encounters a problem or fails, a passive instance can take over.

⁴http://publib.boulder.ibm.com/infocenter/iisinfsv/v8r5/index.jsp?topic=%2Fcom.ibm.swg.im.iis.productization.iisinfsv.ha.install.doc%2Ftopics%2Fwsisinst_pln_ha_xmeta_db2cluster.html

The DB2 agent provides many data attributes and variables available for reporting. Detailed descriptions of these attributes can be found in the agent documentation. The Capacity Planning team should define the online reporting they would like to see with the monitoring team. Specific reporting includes “Database workspaces”⁵

Oracle RAC

The Oracle Real Application Cluster (RAC) database is covered by the Tivoli Composite Application Manager Extended Agent for Oracle Database Version 6.3.1.⁶ “The agent monitors Oracle RAC (Real Application Clusters) databases, ASM (Automatic Storage Management) databases, and normal databases, which have a single Oracle instance for each Oracle database”

Below, is an example of an Oracle RAC database:

⁵http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=%2Fcom.ibm.itm.doc_6.1%2Fitmintro.htm

⁶http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=%2Fcom.ibm.itcama.doc_6.2.4%2Fwelcome.htm




Figure 7: Example of an Oracle RAC Database

The Oracle RAC capability is probably the most important of all the database monitoring as it has a shared cache. "Oracle RAC is a cluster database with a shared cache architecture that overcomes the limitations of traditional shared-

nothing and shared-disk approaches to provide a highly scalable and available database solution for all your business applications.”⁷ This agent provides extensive coverage for RAC.

Oracle RAC performance monitoring

Oracle Real Application Clusters (RAC) enables a single database to run across a cluster of servers. The most important services in RAC are GCS (Global Cache Services) and GES (Global Enqueue Service). The Oracle Database Extended agent provides the following features to monitor those two services:

GCS monitoring

- GCS Memory: the memory that currently used by GCS
- GCS Current Block: the received time and served time of the current block
- GCS Block Lost: the percentage of block lost during the transfer
- GCS CR Latency: the received time and served time of the CR block
- GCS Waits: the waiting sessions that are related to GCS
- Cache Fusion: the ratio of cache fusion

GES monitoring

- GES Enqueue Statistics: the contents statistics in GES
- GES Latency: the latency statistics in GES
- GES Lock: current locks in GES
- GES Memory: the memory that currently used by GES
- Messaging Statistics: the time to process messages in GCS and GES

Documentation shows that the Oracle Extended Agent data can be viewed through the TEP. TCR can be utilized as well for batch reporting. Specific reports that are available within the TEP workspaces include:

⁷ <http://www.oracle.com/technetwork/database/clustering/overview/ds-rac-11gr2-134513.pdf>

- **Cache Fusion Report** - Provides detailed information about cache fusion.
- **GCS Current Block Report** - Provides statistics of time that is consumed for the current block.
- **GCS Memory Report** - Provides detailed information about different parts of memory that are allocated for the GCS.
- **Messaging Statistics Report** - Provides detailed information about messages in the GCS and GES.

Examples of these workspaces are found below:

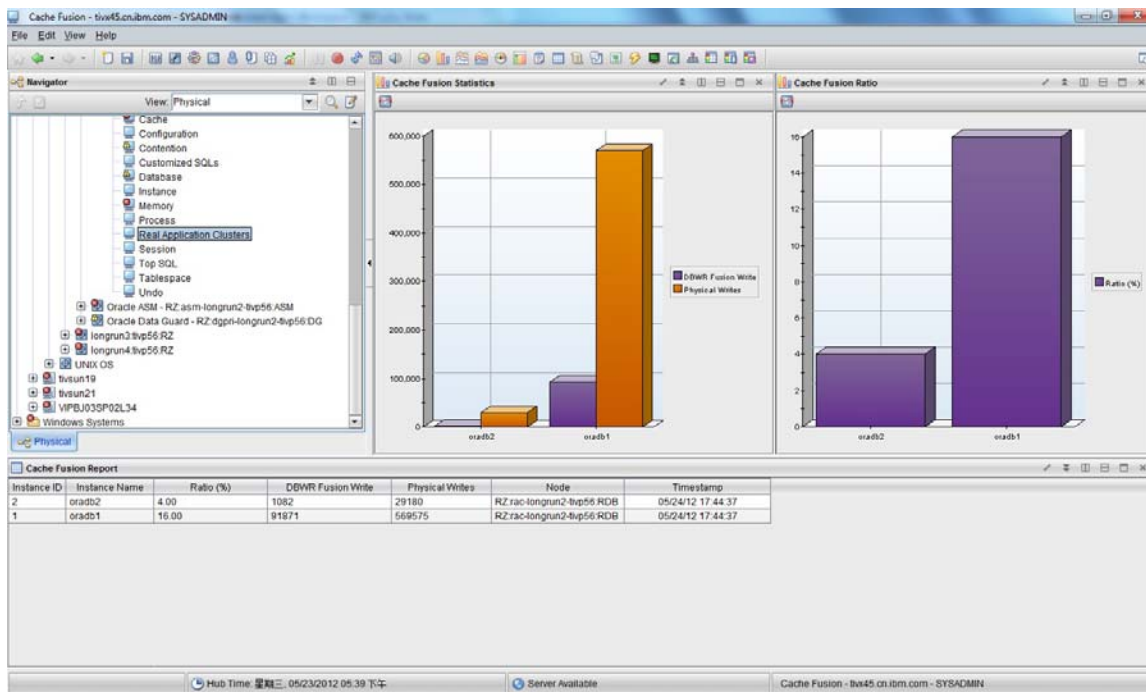


Figure 8 – Cache Fusion Report

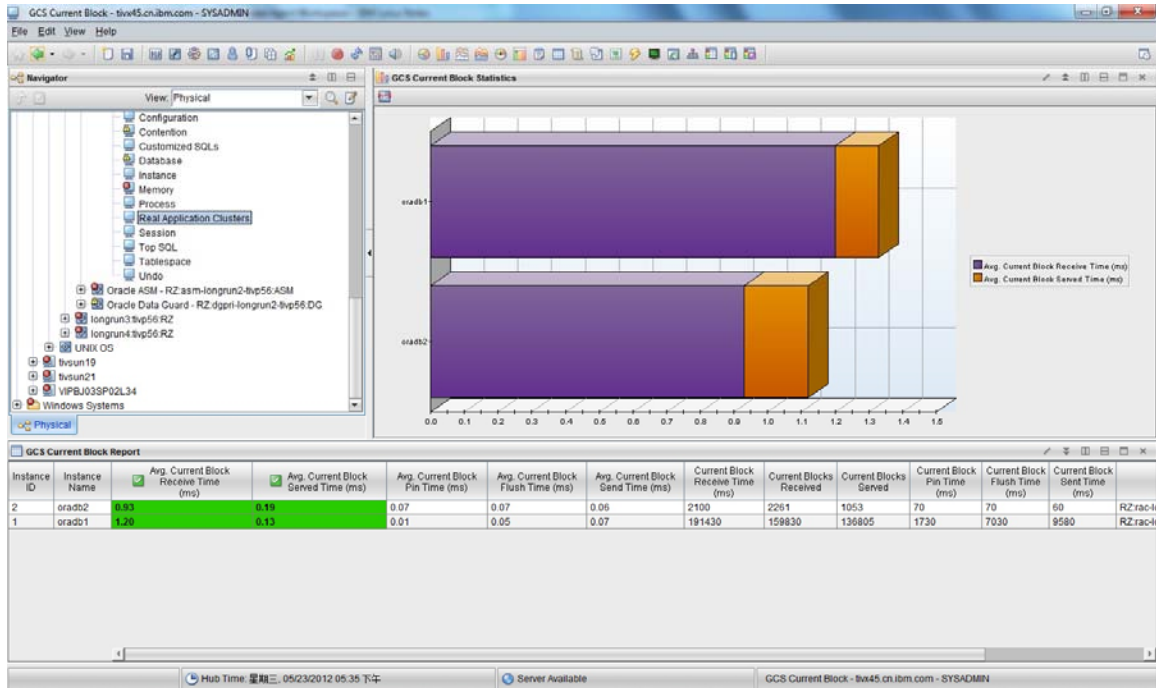


Figure 9 – GCS Current Block Report

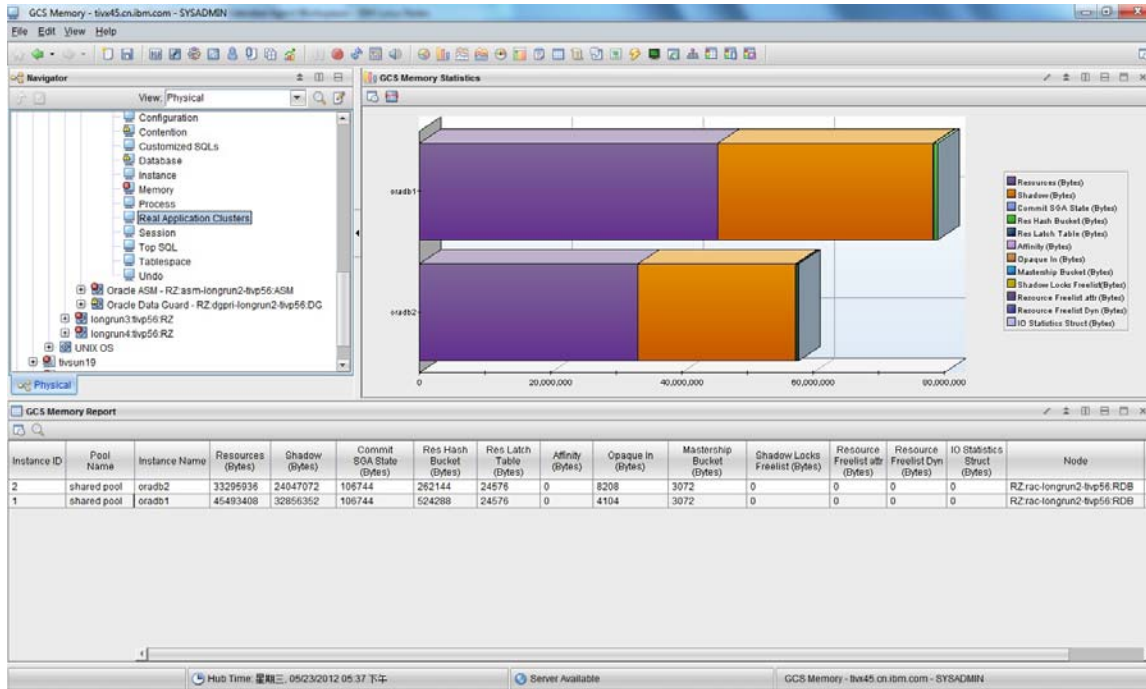


Figure 10 – GCS Memory Report

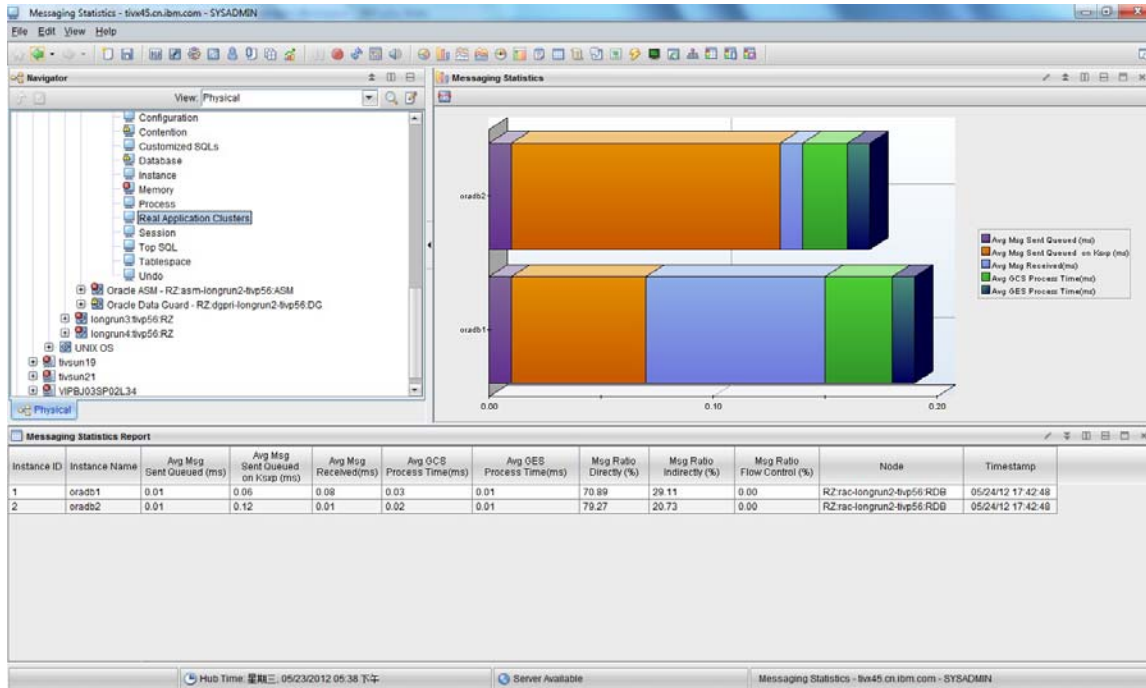


Figure 11 – Messaging Statistics Report

Two sample TCR reports are also provided in the ITCAM for Applications Reports for Tivoli Common Reporting package which can be found in the IBM Integrated Service Management Library⁸. The reports included in this package are:

- *Top n RDB Resource Limitation Report* - This report shows the top resource utilization for the top 10 (or n) systems.

⁸ https://www14.software.ibm.com/webapp/iwm/web/reg/download.do?source=tivopal&S_PKG=1TW10CP1N&lang=en_US&cp=UTF-8

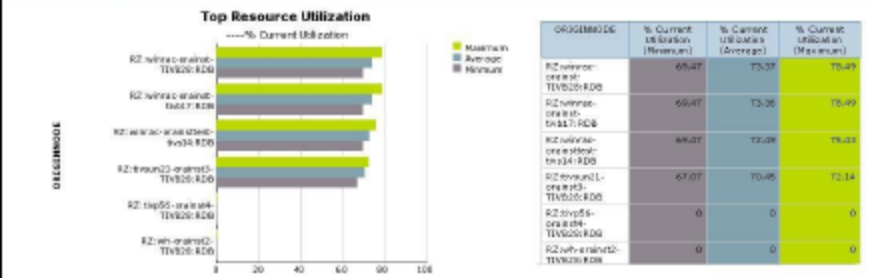
Top n RDB Resource Limitation Report

Oracle Database Extended Agent

Date filter: Last 31 days Vacation Period: All days
Start Date: Dec 5, 2010 12:00:00 AM End Date: Jan 4, 2011 11:59:59 PM
Shift Period: All Shifts Summarization Type: Daily
Resource Name: gcs_resources Number of Systems: 6
Origincode: RZ:hwac21-ora115-TIV020:RDB

The report shows the % current utilization of top RDB resource utilization and inferring agent detail. The upper chart show the information of current time, the lower chart showing the data from history collection and both the detail info of your choosing agent.

Current Time Display



History Collection & Agent Detail

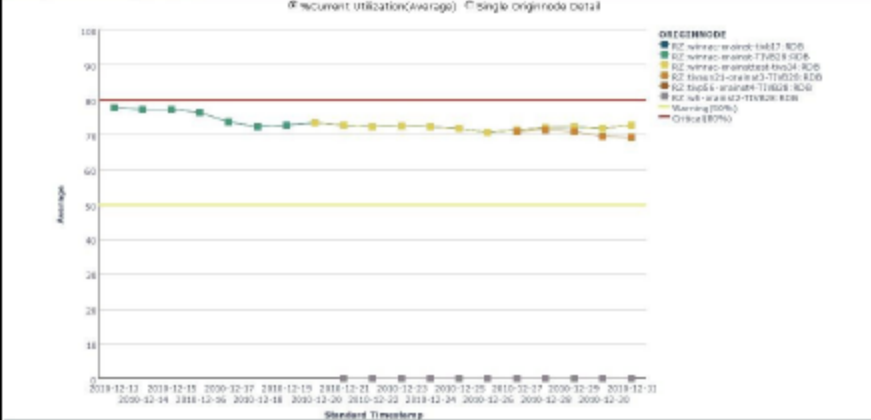


Figure 12: Top n RDB Resource Limitation Report

- RDB Non-Temporary Tablespace Usage Report** - This report shows the % Free To Allocated and % Free To Maximum detail of non-temporary tablespace usage which you selected monitoring database.

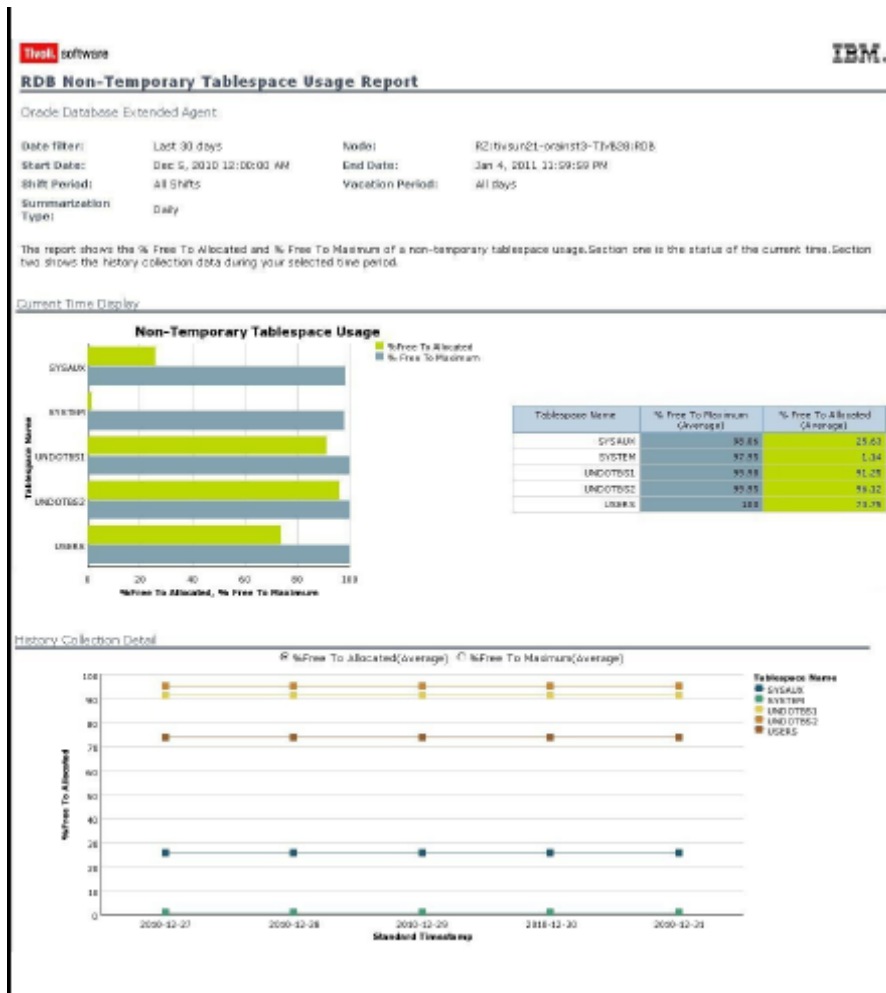


Figure 13: RDB Non-Temporary Tablespace Usage Report

There are many other data elements available that can be used for custom reporting, if required.

Sybase

Sybase comes to the table with the Adaptive Server® Enterprise (ASE) 15 Cluster Edition⁹. This utilizes a shared-disk clustering implementation based on virtualization. Specifically, “ASE Cluster Edition uses a shared-disk architecture designed to support up to 32 ASE server instances that simultaneously access a Storage Area Network. The resources of each ASE server instance in the cluster are virtualized and managed within logical clusters.

A logical cluster is a grouping of ASE server instances, each having dedicated resources from the physical cluster along with unique routing, load balancing and failover rules. Workloads or applications are assigned to a logical cluster, virtualizing them across the servers in the cluster. Administrators create and manage logical clusters using the Workload Manager.”

⁹ <http://www.sybase.com/products/databasemanagement/adaptiveserverenterprise/ase-cluster-edition>

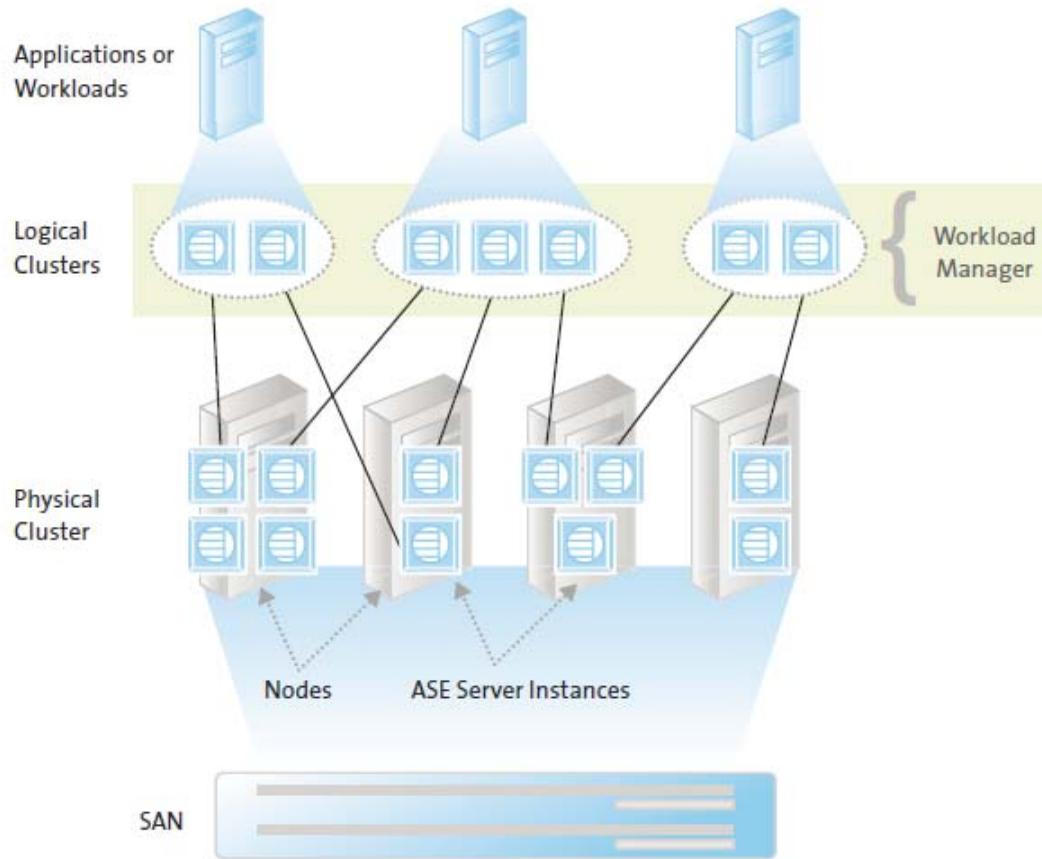


Figure 14: Sybase Physical and Logical Clusters

Here, you can see that there is a many to many relationship between the applications and workloads and the physical clusters. While the ITCAM agent does not specifically manage to clusters (like DB2 and Oracle RAC), it does manage more than one Sybase application on a server and can handle at least a one-to-one relationship between workloads and physical clusters. Here, for example, the ITCAM agent would omit one physical cluster as there are three applications and four physical clusters.

There is one document specific to the Sybase agent: IBM Tivoli Composite Application Manager Agent for Sybase ASE User's Guide. This user guide provides agent-specific reference and troubleshooting information for configuring and using the Sybase agent.

MS-SQL

Below, please find an example of the MS-SQL cluster layout:

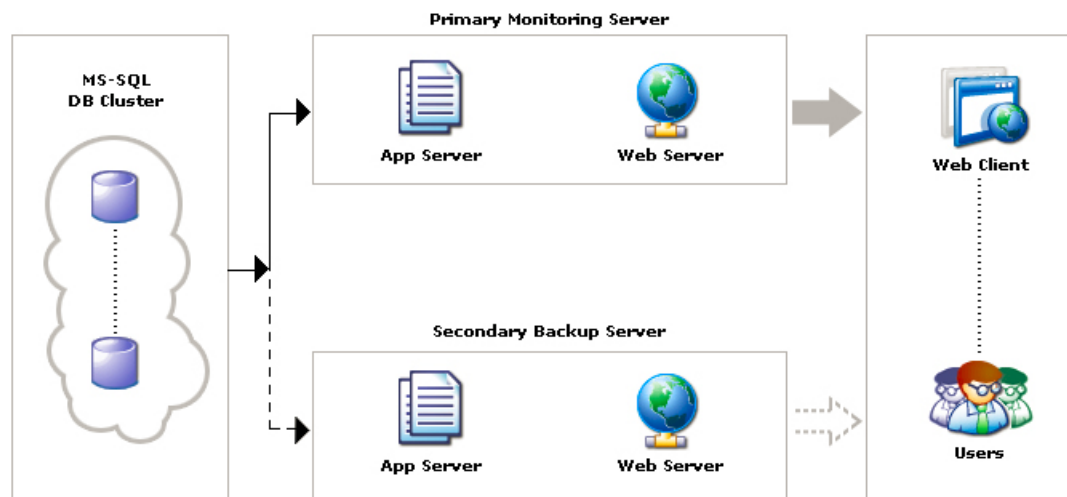


Figure 15: One Possible MS-SQL DB Cluster Layout¹⁰

¹⁰ http://www.manageengine.com/network-monitoring/help/userguide/configure_failover_support.html

IBM Tivoli Composite Application Manager for Microsoft Applications: Microsoft SQL Server Agent Version 6.2.3 covers the MS-SQL tooling for ITCAM¹¹. This documentation shows how to set up MS-SQL in a clustered environment. Tivoli provides support for a cluster environment using Microsoft SQL Server 2005 or Microsoft SQL Server 2008 Enterprise Edition. It appears that the clusters are monitored separately if the agent is set up correctly.

There are several reports available, including:

Category 1: Performance reports

(11 available, see references for details)

Category 2: Utilization reports

CPU Utilization report - Displays the percentage of the CPU time used by the SQL Server processes and the percentage of CPU time used by all processes.

Space Utilization by Database report - Displays the free space (in MB), percentage of free space, the amount of space used as percentage of the total space, the percentage of free space in the transaction log, and the percentage of transaction log that contains the database.

Category 3: Availability reports

(7 reports available, reports shown below are pertinent)

Free Space Availability report - Displays the lowest percentage of free space in the data segment on a device allocated to a database and the lowest percentage of free space in the log segment on a device allocated to a transaction log.

File Group Details report - Displays the amount of space used in the file group as a percentage of the total space allocated for a selected database.

Device Details report - Displays the free space (in MB) on a device, percentage of free space on the device, and the total space (in MB) on a device.

Mirroring Status report - Displays the state of synchronization and mirroring state of a database on a selected server.

¹¹http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=%2Fcom.ibm.itcama.doc_6.2.4%2Fwelcome.htm

Top Databases by Mirroring Status report - Displays the state of synchronization and the state of mirroring of a database on top servers.

Informix

There are many types of Informix applications, ranging from small, medium and large applications and ranging across many different operating systems.¹²

Informix can be used against an Informix database server or against other types of databases (i.e., DB2 or any ODBC of JDBC DB). Its main strength is the application development for optimizing OLTP transactions.

¹² <http://www.redbooks.ibm.com/redbooks/pdfs/sg247884.pdf>

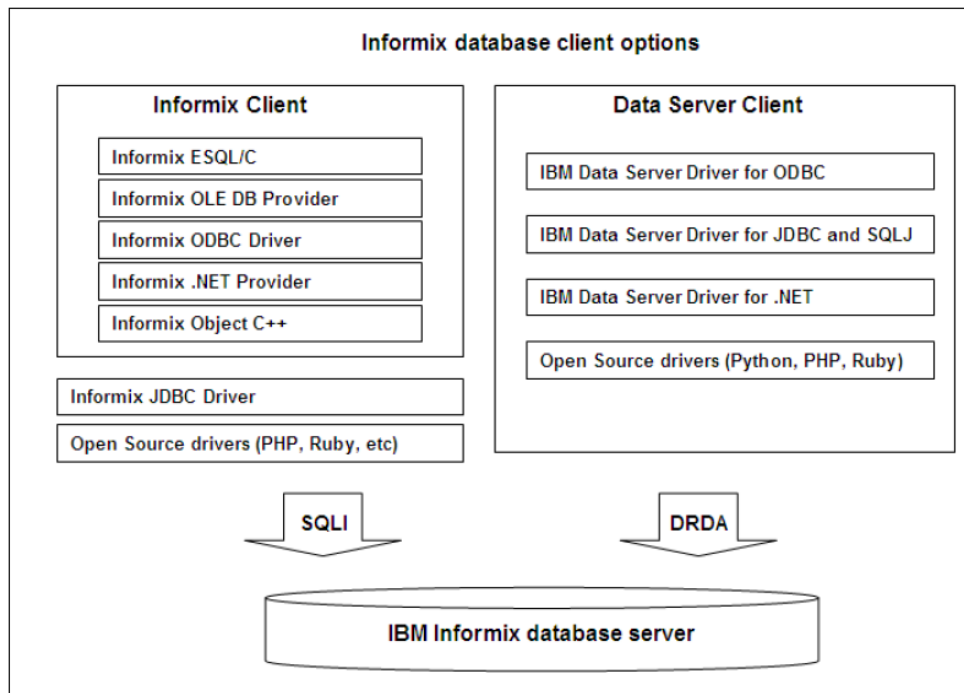


Figure 12: Available Informix Client products for Application Development

Within the application system, the following drivers are available for use. This shows the variety of databases that Informix can run against.

The following IBM Data Server Drivers are available:

- IBM Data Server Driver for JDBC and SQLJ
- IBM Data Server Driver for ODBC and CLI
- IBM Data Server Driver for .NET
- IBM Data Server open source drivers such as PHP or Ruby

Therefore, it is important to see if the Informix systems are using a native (although probably customized) Informix DB or are going after any of many other standard database types.

Informix ITCAM Older Solution

There is an Informix Tivoli monitoring solution that exists and which is highly customizable. “IBM Tivoli Monitoring for Databases: Informix resource models capture and return information, such as database status and server availability, about a resource or software application in the Tivoli management environment. You can change many of the settings for a resource model to customize it for your monitoring needs.”¹³

Various versions of Informix have differing clustering configurations. Looking at the latest version 11.7, “Heterogeneous flexibility leverages existing resource investments and enables scaling beyond a single cluster. Different Informix clusters can be brought together enabling applications to see an entire cluster as a single node. Workloads can be managed and spread across different nodes within this heterogeneous environment. Administrators can dynamically add or drop Informix server nodes to meet changing workload requirements.”¹⁴

There is one cluster environment located in Windows. “Cluster utility (Windows only) and this component is used to provide Microsoft Cluster Server support in Informix.”¹⁵

¹³ http://publib.boulder.ibm.com/tividd/td/ITMD/SC23-4728-00/en_US/PDF/SC23-4728-00.pdf

¹⁴ <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&supplier=897&letternum=ENUS210-367>

¹⁵ <http://www.redbooks.ibm.com/redbooks/pdfs/sg247666.pdf>

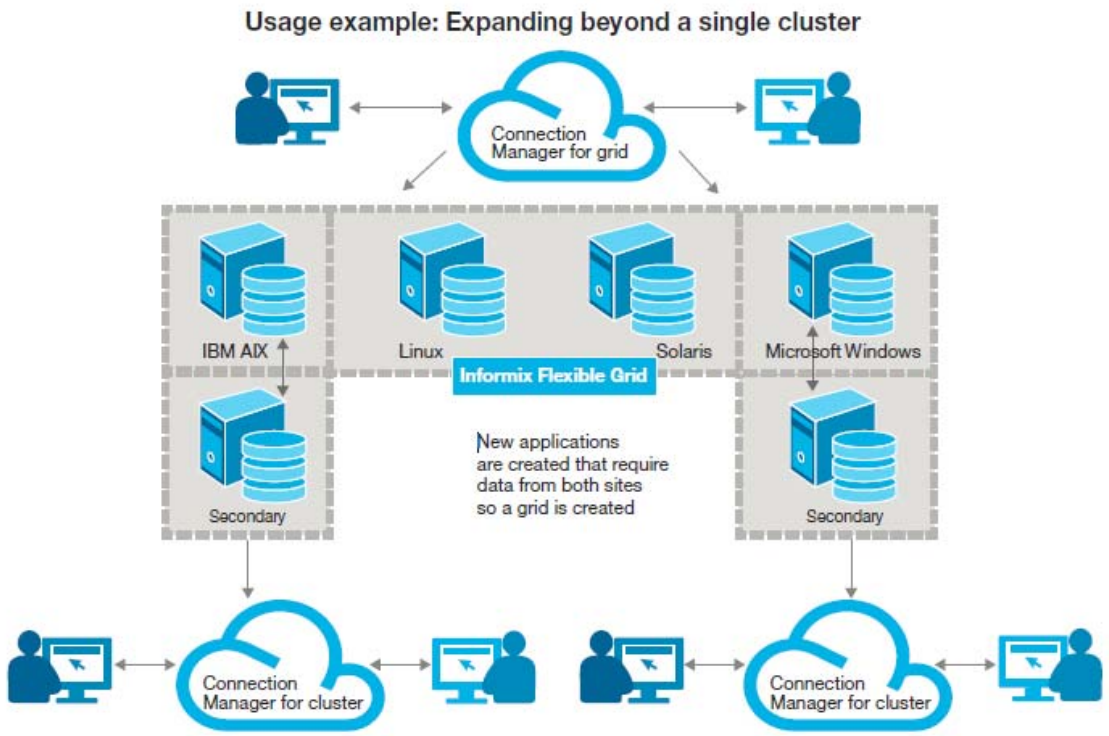


Figure 16: Informix v11.7 Heterogenous Environments

GSMA Newer Informix Scripting Element

GSMA has a newer custom Informix scripting element that was built off the older OPAL solution (described above).¹⁶ This element covers the following solutions:

¹⁶ <http://gsma.lagaude.ibm.com/articles/ua/1065>

Standard Base monitors for Informix					
Monitored Resource	General Description	Status	Event Severity	Polling Interval (in seconds)	Occurrence
IBM Informix Server Status	Informix Server in a state other than online	Not 'OnLine'	FATAL	60	1
IBM Informix checkpoint duration	Informix Checkpoint duration too long	Checkpoint duration exceeds 5s	CRITICAL	300	1
IBM Informix Logical Log Backup Duration	No Logical Log Backup since too long	No Logical Log Backup since 1h	CRITICAL	1800	2
IBM Informix Table Extents Number	Number of Table extents too high	Number of Table Extents exceeds 90	CRITICAL	3600	1
IBM Informix Database Free Space	Database Free Space too low	Database Free Space under 10%	CRITICAL	300	1
Optional recommended monitors: Common Agent standard situations for Informix log monitoring					
Logfile Monitor FATAL Entries	Issues a FATAL alert if pattern is found in a logfile	rc==17 and severity=='FATAL'	FATAL	N/A	1
Logfile Monitor WARNING Entries	Issues a WARNING alert if pattern is found in a logfile	rc==17 and severity=='WARNING'	WARNING	N/A	1

Figure 17: Currently Supported Informix Base Monitoring Solution from GSMA

Note that these monitors do not specifically cover clustered databases. Therefore, a section has been added for custom CLI scripting to cover any clustered database use for Informix in the overall design document.

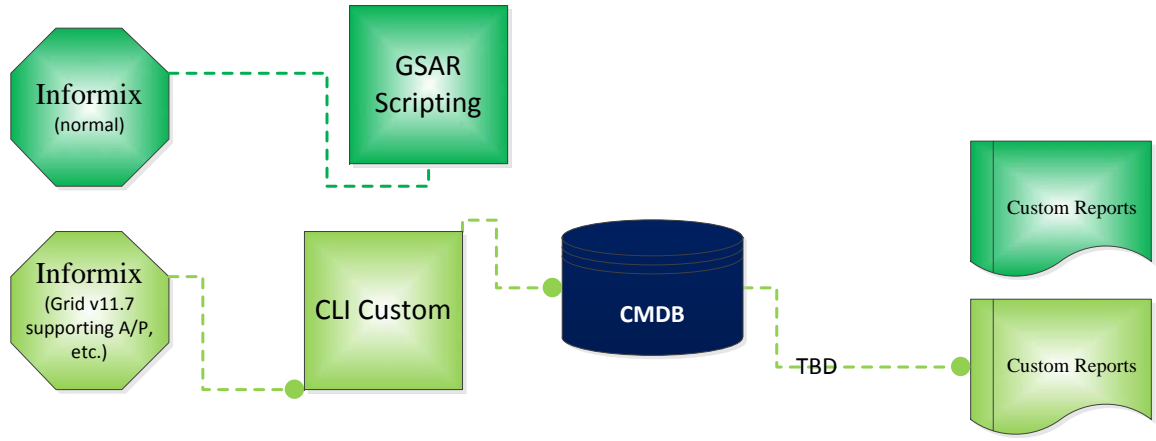


Figure 18: Informix Blended Solution


CONCLUSION

This paper has shown architecture for utilizing Tivoli's Composite Application Manager (ITCAM) to provide reporting for many clustered databases and differentiates between those architectures which are supported (such as Oracle RAC and MS-SQL) and which require custom tooling (such as Informix). While standard reporting is then identified for each supported database system, the actual implementation of the reporting outlined in this paper was constrained to disk utilization and tablespaces, to fulfill an SLA requirement. .

Appendix A: TPC Reports (non-Clustered DB)


History Chart: % of Space Used

Created by: tpcadmin Printed: Mon Feb 13 16:06:32 PST 2012



History Chart: Free Space

Created by: tpcadmin Printed: Mon Feb 13 16:04:48 PST 2012



Tablespaces: By DBMS Type

Created by: tpcadmin Printed: 2/9/12 12:18 PM

RDBMS Type	Computer Name	Instance	RDBMS Version	Port	Home Directory	
UDB	sssatest	DB2	SQL09070	50,000	C:\Program Files\IBM\SQLLIB\	
Logonid	Probe Status	Last Start Time	Last Probe Time	Tablespace Count	File Count	Log File Count
db2admin	Succeeded	Feb 9, 2012 10:15:01 AM	Feb 9, 2012 11:17:12 AM	10	10	8
File Free Space	Total Size	File Capacity				
75.87 MB	392.61 MB	236.29 MB				

TPCUser.UDB – Instance Storage By Computer

Created by: tpcadmin Printed: 2/9/12 12:01 PM

Computer	Total	Container	Container	Log File	Tablespace	Container	Log File
	Size	Capacity	Free space	Capacity	Count	Count	Count
sssatest	392.61 MB	236.29 MB	75.87 MB	156.31 MB	10	10	8