
Crystal Enterprise Scalability and Sizing Benchmark

- ▶ *Crystal Enterprise 10 and IBM AIX 5.2*
(Tests Conducted at IBM Solution Partner Center—San Mateo)

Author: Keith Moon

Contributors: Cortney Claiborne, Kevin Crook, Davythe Dicochea, Erin O'Malley, Derek Stobbart, James Thomas

Distribution under NDA

Contents

Executive Summary: Crystal Enterprise 10 Scalability on IBM AIX	1
Benchmark Test Overview	1
Test Configuration Overview	1
Results Summary	2
Introduction	3
Test Concept	4
User Actions	4
Report Design	4
System Environment	4
Resources and Resource Usage	4
Reliability	4
Benchmark High-Level Design	5
Test Procedure	6
Test Results	7
CPU Scalability Measurements 12–24–48 (Peak Load)	7
<i>Comparative Scaling Test Sequence</i>	7
<i>Active Users Tested and Response Times</i>	9
<i>Requests per Second</i>	9
<i>Throughput</i>	10
<i>Volume of Requests by Type</i>	10
<i>CPU Utilization</i>	11
Conclusion	12
Appendix 1: Scalability Primer	13
Scalability	13
Performance	13
Speed Spot	14
Sweet Spot	14
Peak	14
Reliability	14
Functionality and Usability	14
Appendix 2: IBM Total Storage Proven Certification	15

Appendix 3: Test Environment	16
Load Harness	16
Weighted Script Mix	16
Script Parameters	17
Think Times	17
User Definitions	17
Load Schedule Definition	18
Content Check (Error Checking)	18
Working Set	19
Appendix 4: Reports	20
Design	20
Appendix 5: Report Data Source	22
Appendix 6: Software Environment	23
Crystal Enterprise 10 Configuration and Tests	23
WebSphere	23
DB2	23
Appendix 7: Topology	24
Appendix 8: Hardware Environment	25

Executive Summary: Crystal Enterprise 10 Scalability on IBM AIX

Business Objects conducts extensive scalability and performance testing on our products to get a thorough understanding of how the software performs in actual customer deployments. In addition to this testing, comprehensive performance benchmark testing is done with third parties to verify internal results and to benchmark our software based on a more extensive set of real-world implementation scenarios.

In February 2004, at the IBM Partner Solution Center in San Mateo, California, benchmark tests were performed on Crystal Enterprise™ 10 (CE 10) on IBM AIX hardware to provide a thorough understanding of system performance, reliability, and stability under various user loads, for different report types, and in different configurations. Business Objects uses the results to help customers understand the overall scalability of the system to help plan for actual deployment of the software. While benchmark tests are often used to compare the scalability of competitive systems, Business Objects believes that providing test results for real-world customer deployment should be the ultimate goal of performance testing. Therefore, we designed test scripts and configurations that map to actual deployment scenarios with real report samples and data sources.

Benchmark Test Overview

The tests proved that Crystal Enterprise continues to provide outstanding report-processing performance in terms of throughput of data and viewing response times.

The tests performed include the scalability, performance, and reliability within simulated real world conditions. All tests were performed on a 32 way p690 (LPAR into four 8ways) to test the ability of the software to use additional processing power efficiently. The key measure of efficiency is linear scalability or, simply, the ability of the software to perform consistently as hardware and processing power is added. Due to latency with internal system communication protocols, perfect scalability (1:1) is impossible to achieve. However, many software systems experience non-linear scalability resulting in degradation of performance with more resources. A system cannot be considered scalable if there is significant degradation.

The test scripts used were designed around real-world scenarios. As such, they must include think time, or the amount of time a user actually looks at a report once it is processed. This is an important factor in customer testing because it means that the report must maintain state within the system. And this is important because refreshing the report does not require resubmitting the query to the database, which would introduce unnecessary load on the database. Tests with no think time built in are not real-world and are deceptive in terms of performance testing.

Test Configuration Overview

- Crystal Enterprise 10 installed on IBM AIX 5.2 machine with 4, 8, 16, and 32 CPU configuration enabled.
- Five different report scenarios from a single page report to a multi-page drill report.
- 500 to 4,000 active users. This translates to roughly 40,000 concurrent users, respectively, and 400,000 named users (assuming a 10 to 1 ratio of active users to concurrent users viewing or interacting with the report at any one time).

Results Summary

- Achieved a consistent, positive, linear scalability of over .92 (for every additional processor added, we achieved a greater than 92% load increase). Only the best software products in the world can boast this claim.
- Impressive 14.3 MB per second throughput with 32 CPUs and 4000 concurrent users.
- Highly predictable, quick, and efficient responsiveness across all tests. Sub 2, 4, and 6 seconds response times (for 90th percentile) including live requests to database.
- Proven system fault tolerance under load.

Introduction

The benchmark tests conducted at the IBM Partner Solution Center in San Mateo, California allowed Business Objects to test Crystal Enterprise 10 for the purposes reporting levels of scalability, performance, and reliability within a simulated real-world condition.

The benchmark test (particularly the workflows, reports, and database) was not designed or skewed to show an incredible number of concurrent virtual user connections. The variations of functionality, complexity of reports, weight of requests per second was determined in order to provide a realistic concentration of load. Increasing or decreasing the complexity of the design would decrease or increase the capacity of concurrent virtual users.

Based on the results obtained from this particular test design, this paper will serve as a valuable sizing tool for system analysts to estimate and project potential environment requirements over a wide resource spectrum.

Test Concept

This benchmark was inspired and conceptualized based on our customer feedback as to how they implement and use Crystal Enterprise. Test design was also influenced directly based on the results of Customer Advisory Meetings, Surveys, Proof of Concepts, etc. The external benchmarks for Business Objects are designed as a hybrid test environment that reflects a common ground of integrated use. The sequence of tests was chosen to report behavior over a spectrum of user load and resources while maintaining the same integrated test scenario.

User Actions (Functionality and Mix)

Test scenarios (workflows) were designed to emulate common user interaction and requests to the system using realistic variable think times. The functionality includes logging on to the system, navigating folders, selecting reports, viewing reports containing saved data, viewing reports requesting live data, navigating from page to page within a report, drilling down through groups, scheduling reports, exporting to PDF or Excel, and logging off.

Think times are used within the scripts to help create a more accurate model of user load and system resource use.

Report Design

The tests use a mixture of reports that range from smaller data sets up to larger data sets and simple design up to high complexity. For example, a one-page summary report with 1,000 records to 600-page report with 100,000 records. Reports are comprised of formulas, grouping, charts, images, parameters, and drill-down navigation. See Appendix for sample reports.

System Environment

The Crystal Enterprise environment (working set) consists of a realistic total number of objects (users, reports, folders, groups, servers). See Appendix for system configuration.

Resources and Resource Usage

Acceptable virtual user load levels consider available resources and assure that resource usage (CPU and memory) and performance (response times, throughput) falls well within normal acceptable limits. The test sequence was designed to test loads that range from highest performance (speed spot) to peak level usage.

Reliability

System configuration was designed for an equal balance of performance, scalability, and reliability. It was not configured to merely prove performance as many other benchmarks do.

Performance: Enough services (e.g., Page Server, Cache Server, etc.) were made available to support total user load in order to maintain fast response times.

Scalability: Services were distributed to assure the highest degree of load balancing and to support system growth.

Reliability: Services were distributed to assure both software and hardware fault tolerance. If a service were to fail, another service would be available to support requests. If a machine were to fail, another machine would be available with all services to support requests.

Benchmark High-Level Design

This benchmark consisted of a test series that cross the spectrum of low to high virtual user loads and from lower to higher number of CPUs and memory. The specific series was selected to concurrently demonstrate scalability capabilities, performance under various loads vs. hardware combinations, and reliability.

All virtual users are always active users (users who are using up RAM and CPU cycles by making requests for navigating, viewing, scheduling, etc).

4 (4 x 1CPU)	8 (4 x 2CPU)	16 (4 x 4CPU)	32 (4 x 8CPU)
2CPU WAS	4CPU WAS	8CPU WAS	16CPU WAS
			4000 virtual users
			3200 virtual users
		2000 virtual users	2000 virtual users
		1600 virtual users	1600 virtual users
	1000 virtual users	1000 virtual users	1000 virtual users
	800 virtual users	800 virtual users	800 virtual users
500/600 virtual users	600 virtual users	600 virtual users	
400 virtual users	400 virtual users	400 virtual users	
200 virtual users	200 virtual users		
100 virtual users			

Test Procedure

- 1) Working set at baseline (delete any previously scheduled report instances from previous test)
- 2) Recycle test machines
- 3) Initiate performance monitoring
- 4) Start prescheduled recurring instances [Large Report (100K–600 Page) @ 5 minute intervals, Schedule to Export (10,000K–50 Page) @ alternate 5 minute interval]
- 5) Start load harness suite (ramp up 100 user every 13 seconds)
- 6) Run at peak load for 30 minutes (full virtual user load)
- 7) Ramp down (gradual shutdown or immediate shutdown acceptable)
- 8) Results analysis
- 9) Repeat pass (Steps 1 to 7)
- 10) Results analysis

Test Results

- The results from the benchmark tests demonstrate excellent linear scalability and predictability (ranging from 100 users to 4,000 users) across hardware configurations and resources ranging from a combined total of six CPUs to a combined total of 48 CPUs.
- The results indicate that when using a load model of integrated functionality and an environment subject to increasing levels of concurrent active users and increasing resources, performance levels can be predicted and maintained linearly.
- The results show that all request types (logon requests, view requests, navigation requests, drill downs, and schedules) maintain performance as the system scales.
- The results show that under a reasonable load, acceptable response times (under six seconds) are maintained for 90% of total concurrent active virtual users as a system scales upwards and outwards.
- The results show an impressive degree of throughput (14.3 MB per second) managed by the Application Server and supported at every tier.

CPU Scalability Measurements 12–24–48 (Peak Load)

The following results represent peak loads across all hardware configurations. The four tests in this sequence demonstrate that as user load is increased proportionately to hardware resources, performance, and reliability, functionality is maintained.

The scalability results were achieved with and by including the following configuration aspects (configuration details):

- Configurable number of processes as required to support load
- Configurable multithreaded tuning as required to support load
- Configurable redundant distributed architecture to support performance and reliability
- Configurable multi-user caching (saved data reports only for this test—on demand reports always query the database)

Comparative Scaling Test Sequence

Test 1—500 active users loaded on total of 6 CPUs

Test 2—1,000 active users loaded on total of 12 CPUs

Test 3—2,000 active users loaded on total of 24 CPUs

Test 4—4,000 active users loaded on total of 48 CPUs

	Test 1 4 (4 x 1CPU) 2CPU WAS 500 virtual users	Test 2 8 (4 x 2CPU) 4CPU WAS 1000 virtual users	Test 3 16 (4 x 4CPU) 8CPU WAS 2000 virtual users	Test 4 32 (4 x 8CPU) 16CPU WAS 4000 virtual users
Load	Peak Load	Peak Load	Peak Load	Peak Load
Avg. Throughput (bytes/second)	1,932,874	3,724,585	7,477,570	14,306,245
Scalability Factor		1.927	3.869	7.402
Scalability %		96.3%	96.7%	92.5%
Total Hits	87,029	171,072	347,572	669,395
Average Hits per Second	48.349	95.04	193.096	371.886
Total Throughput (bytes)	3,479,173,379	6,704,253,274	13,459,625,274	25,751,241,885
Response Times (sec) Median				
Live Data	5.0	4.7	4.0	4.1
Saved Data	1.3	1.8	1.0	0.6
Logon	1.7	1.8	0.2	0.7
CMS Query	1.1	1.0	0.4	0.2
Schedule	0.9	1.0	0.6	1.2
Drill Down	1.5	2.3	1.4	1.3
CPU % of Web Application Server	87%	91%	88%	90%
CPU % of Crystal Enterprise Report Servers	82%	84%	81%	84%
Response Success Rate*	100%	99.8%	100%	99.9%
System Stability (server uptime)	100%	100%	100%	100%

**Response success rate based on percentage of requests returned as successful response. Any failed responses cause the virtual user to exit the script and immediately choose a new script based on current weighting.*

Active Users Tested and Response Times

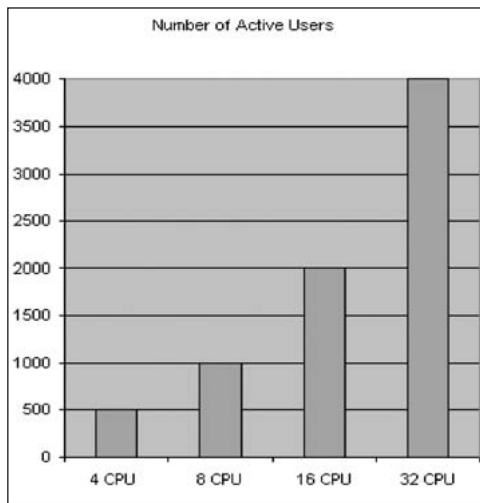


Figure 1: shows the number of users increasing with additional CPUs.

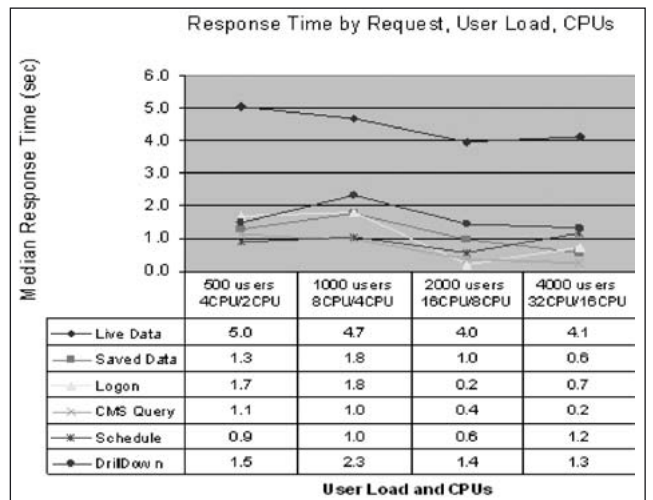


Figure 2: shows the performance levels being maintained at the same high standard across several configurations.

Requests per Second

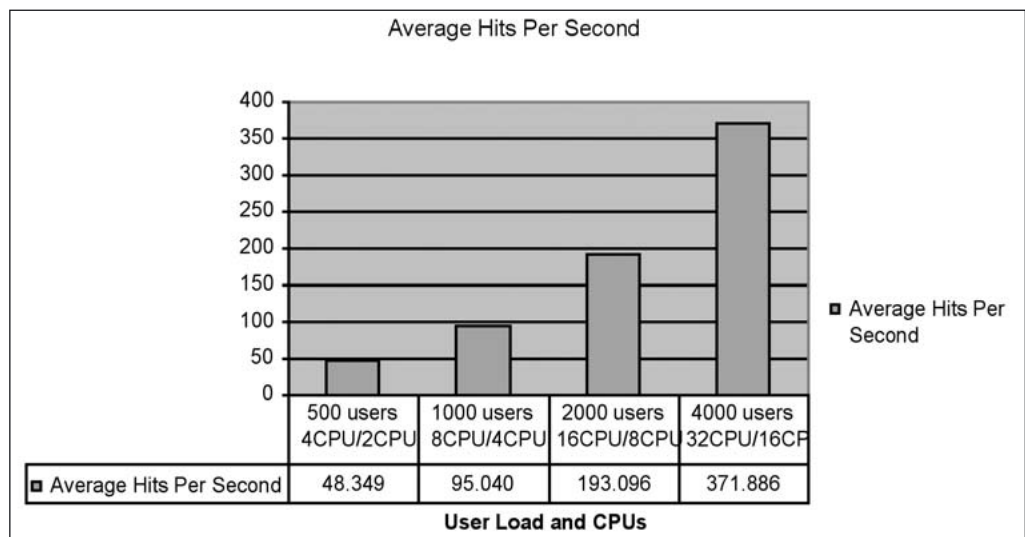


Figure 3: shows that as the user load and number of CPUs increases, the number of hits per second increased at the same rate.

Throughput

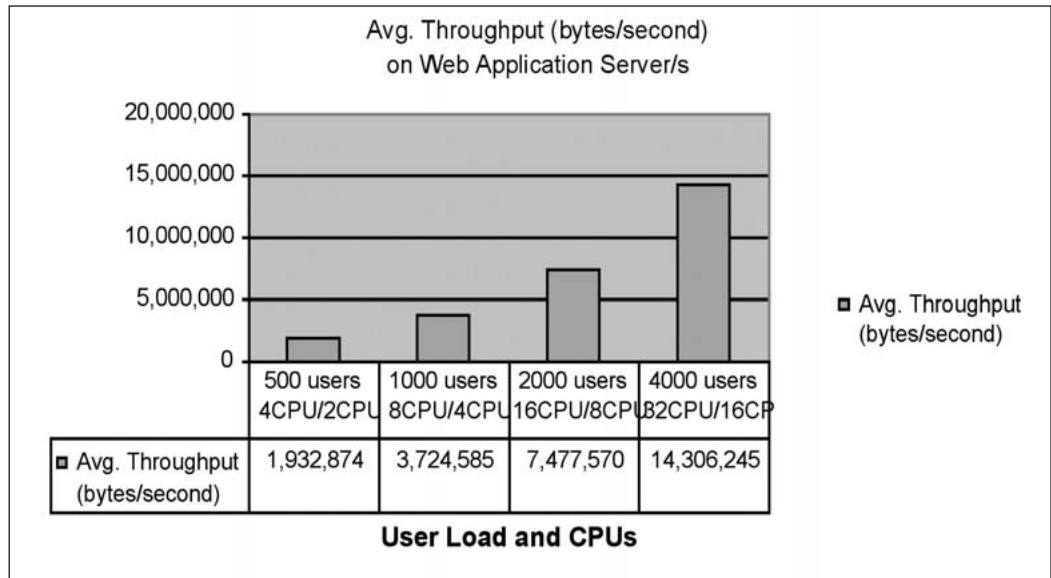


Figure 4: shows that as the data load increases in line with the user load and CPU, the throughput increases at the same level.

Volume of Requests by Type




Figure 5: shows the number of requests for each operation increased in line with the increase in CPU resource and user load.

CPU Utilization

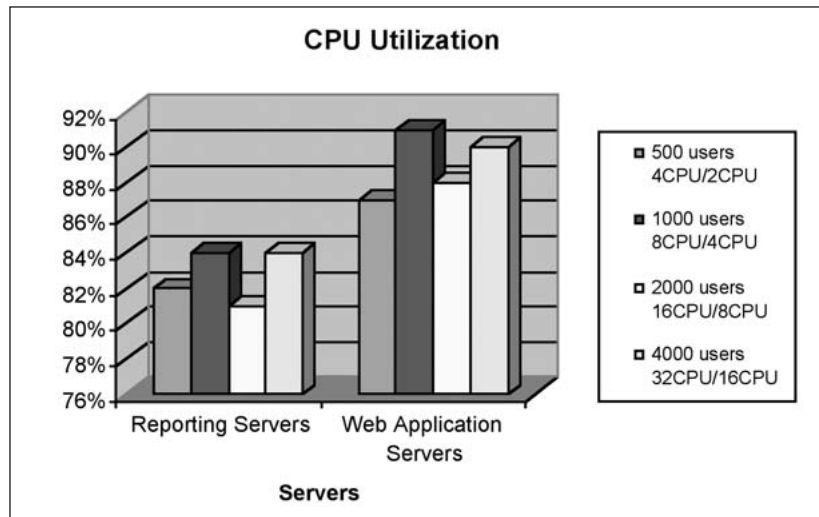


Figure 6: shows that as the measures were taken at each level, the system use remained at an acceptable working level and was consistent across all the tests.

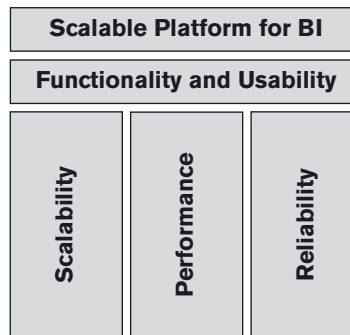
Conclusion

Crystal Enterprise 10 continues to provide a high performance, scalable, and reliable business intelligence (BI) and reporting platform for a wide variety of customer scenarios. These tests prove that Crystal Enterprise delivers consistent linear scalability across broad deployment scenarios for virtually any size and complexity of deployment.

Business Objects encourages customers to review the benchmark tests in conjunction with the professional services provided by our company or partners when designing, configuring, and deploying their BI systems.

Appendix 1: Scalability Primer

Real-world scalability depends upon a balance of reliability, scalability, and performance that supports available functionality and usability:



Scalability

The goal of this benchmark was to demonstrate Crystal Enterprise 10 and its ability to scale as additional transactions or users—as well as an equal amount of resource parts—are added while maintaining performance (e.g., response time), usability, and functionality.

The ability for the whole system to scale will directly determine and help to predict future resource costs as system use increases.

Performance

The benchmark test series was designed within a spectrum that shows what kind of performance can be expected when users (with realistic usage) and resources (CPUs, RAM) are added.

The following generic chart illustrates the concepts used in determining test sequence design:

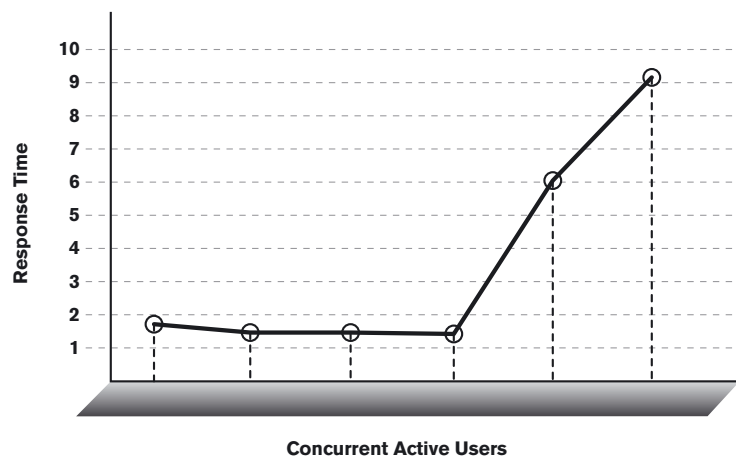


Figure 7: Performance Region. This chart shows how response time varies as user load is increased on fixed hardware (4 CPUs).

Speed Spot

The speed spot in the graph is the point where the system begins to perform under any degree of stress. Prior to the speed spot, the system is being under utilized (i.e., excessive hardware ratio to load or 10–15% CPU utilization).

Sweet Spot

The sweet spot is the area of optimal performance. The system is under stress, properly utilizing available resources and maintaining an acceptable level of performance while still leaving room for any normal degradation if extra load is required (i.e., 50–75% CPU utilization).

Peak

The peak is the area where system resources are being stressed to the point where there is no further room for degradation (100% resource utilization). This peak is an area that should be avoided in every day system usage. Before this point is reached, additional hardware resources should be considered. Benchmark results demonstrate through upward and outward linear scalability. Adding new resources will allow for predictable performance gains.

The performance zone runs between the speed spot and the sweet spot. Depending on throughput or response time requirements, usage patterns, and resource constraints, customers should architect their system to fall within this zone.

Reliability

The benchmark test series was designed within a spectrum that demonstrates flexible, scalable configuration of services that provides both software and hardware redundancy. In all configurations, redundant services provided for software fault tolerance.

During this benchmark, Business Objects requested that the Crystal Enterprise system be certified for IBM's "Total Storage Proven" Certification. The certification was achieved using the same configurations used for scalability and performance testing.

Functionality and Usability

Workflows emulating real-life situations were designed for the test to show, as closely as possible, how the system would behave in a true implementation. Usage routines that included viewing prescheduled reports, viewing of reports accessing data in real time, processing of large reports in the background, and exporting to Excel and PDF were all part of the test to exercise each component of the enterprise reporting framework.

Appendix 2: IBM Total Storage Proven Certification



Testing level achieved:

- Standard: consisting of install, configuration, load, and exercise I/O.
- Tests include failover and recovery support.

Details of the planned tests:

- The overall test displayed the flexibility in set-up of Crystal Enterprise 10 to suit differing deployment scenarios and to match changing volume requirements.
- The test environment simulated a real-world deployment scenario with high availability requirements.
- Failures in the fiber switch and of the individual components in the storage subsystem were simulated to test the ability of Crystal Enterprise 10 to continue providing uninterrupted full service to all connected users.
- During the test, three of the four logical partitions on the p690 Crystal Enterprise 10 system were forcefully removed from the network to simulate server crash, reducing the system from an eight processor installation to a two processor installation.

A description of the actual tests and the results obtained:

- Installation was carried out using standard techniques, documented in the installation manuals. The Crystal Enterprise 10 environment was duplicated for redundancy purposes across all four logical partitions of the p690 server in order to provide failover support for all services.
- 250 virtual users accessing the system were simulated using LoadRunner software. The users were carrying out a variety of activities, such as report scheduling, scheduled report viewing, and on-demand report viewing in order to simulate normal system activity.
- On the introduction of the failures mentioned above, certain virtual users experienced longer than normal response times while the system recovered their connection, but no lasting errors were logged at this time and the system recovered optimum response times shortly after each of the failures.

Business Objects Crystal Enterprise 10 Storage Proven listing:
<http://www.storage.ibm.com/proven/attract.htm>

Appendix 3: Test Environment

Load Harness

LoadRunner Version 7.8—URL-based scripts web (HTTP/HTML) protocol

Weighted Script Mix

Script	Category	Weight
Script 1	Viewing On-Demand	30%
Script 2	Viewing Saved Data	60%
Script 3	Schedule to Format	10%

Script 1	Script 2	Script 3
1. Logon	1. Logon	1. Logon
2. Navigate to Folder	2. Navigate to Folder	2. Navigate to Folder
3. Navigate to SubFolder	3. Navigate to SubFolder	3. Navigate to SubFolder
4. Choose Random Report	4. Navigate to SubFolder	4. Choose Random Report
5. View Report (On-Demand)	5. Choose Random Report	5. Schedule Report
6. Select Random Drilldown Node	6. View Report (Saved Data)	6. View History
7. Select Random Drilldown Node	7. Select Random Page	7. Close Report
8. Select Random Drilldown Node	8. Select Random Page	8. Navigate to Folder
9. Select Random Page	9. Select Random Page	9. Choose Random Report
10. Select Random Page	10. Select Random Page	10. Schedule Report
11. Select Random Page	11. Navigate to Folder	11. View History
12. Close Report	12. Navigate to SubFolder	12. Close Report
13. Logoff	13. Choose Random Report	13. Logoff
	14. View Report (Saved Data)	
	15. Select Random Page	
	16. Select Random Page	
	17. Select Random Page	
	18. Select Random Page	
	19. Close Report	
	20. Logoff	

Script Parameters

username—variable user name
password—user password
cms_name—crystal enterprise CMS name
url_path—root path
servername—webappserver names
reportID—random variable reportid
randompage—random variable report page
Logon_ThinkTime—random variable thinktime for logon
ChooseReport_ThinkTime—random variable thinktime for report select
NavFolder_ThinkTime—random variable thinktime for folder select
NavPage_ThinkTime—random variable thinktime for page select

Think Times

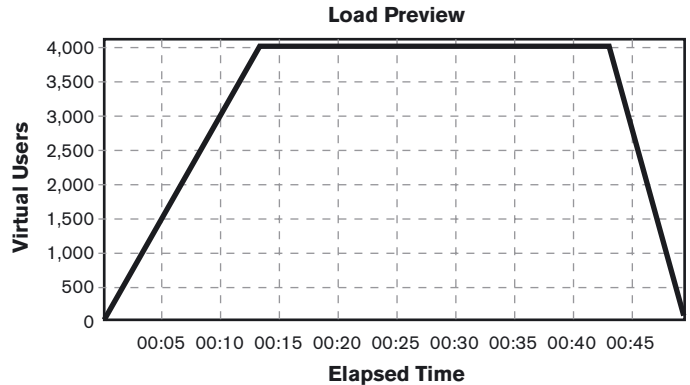
Parameter	Random Number Min (Seconds)	Random Number Max (Seconds)
Logon_ThinkTime	5	10
ChooseReport_ThinkTime	20	32
NavFolder_ThinkTime	2	6
NavPage_ThinkTime	8	12

User Definitions

Total User Population	All users (physical people) who have access to Crystal Enterprise
Concurrent Users	Users who are logged in to Crystal Enterprise. May be using RAM, but might not be using CPU
Active Users	User who are using up RAM and CPU by making requests (navigating, viewing, scheduling, etc.)
Virtual Users	Load Harness Users (includes Active Users, Concurrent Users)

Load Schedule Definition

Virtual users were loaded in a step process until full load. At point of virtual user full load, the test would continue for 30 minutes and then begin to ramp down. For all tests, data is derived from the 30-minute full load period only.



Content Check (Error Checking)

Internet Protocol: ContentCheck Run-time settings

```
<Application Name="CatchError" Enabled="true" StampTime="" StampHost="">
  <Rule Name="Rule_1" FailIfNotFound="false" MatchCase="false"
    UseText="true" Level="error">
    <Text>com.crystaldecisions.report.web.viewer</Text>
  </Rule>
  <Rule Name="Rule_2" FailIfNotFound="false" MatchCase="false"
    UseText="true" Level="error">
    <Text>Account Information Not Recognized</Text>
  </Rule>
  <Rule Name="Rule_3" FailIfNotFound="false" MatchCase="false"
    UseText="true" Level="error">
    <Text>An error occurred at the server</Text>
  </Rule>
  <Rule Name="Rule_4" FailIfNotFound="false" MatchCase="false"
    UseText="true" Level="error">
    <Text>Problem found in</Text>
  </Rule>
  <Rule Name="Rule_5" FailIfNotFound="false" MatchCase="false"
    UseText="true" Level="error">
    <Text>Unsupported</Text>
  </Rule>
  <Rule Name="Rule_6" FailIfNotFound="false" MatchCase="false"
    UseText="true" Level="error">
    <Text>Action canceled</Text>
  </Rule>
</Application>
```

Global Timeout

Any request not satisfied with either a success or failure as defined by the response containing an HTTP response code (e.g., 200, 302, 304, etc.) and content checking will return an error after 120 seconds.

Working Set

-
- 20,000 named users
 - 20,128 report objects
 - 20,010 folders
 - 20,128 report instances (.rpt with saved data)
 - Simple to complex reports containing formulas, groups, cross-tabs, charts
 - Small to larger reports ranging from 1,000 records up to 100,000 records and 10 pages up to 500 pages

Appendix 4: Reports

Design

basicbenchmark.rpt—basic line item report

4 Tables, 17 Fields, 1 Group Field

50 to 600 Page—1,000 to 100,000 records

summariesbenchmark.rpt—medium complexity report containing summaries

4 Tables, 17 fields, 4 Group Field, 40 Summary Fields

1 Page Summary—10,000 Records

featuremixbenchmark.rpt—complex report containing mixture of reporting features (e.g., charts, summaries, groups, parameters, etc.)

4 Tables, 17 fields, 19 Formula Fields, 3 Running Total Fields, 4 Group Fields, 8 Special Fields, 18 Summary Fields, 3 Charts, 1 Cross Tab, 1 Bitmap Image

50 to 2,169 Page—1,000 to 100,000 records

crosstabbenchmark.rpt—report containing cross tabs

4 Tables, 17 Fields, 2 Group Field, 2 Cross Tab, 6 Formula Fields, 1 Special Field

21 Page—10,000 records

textobjectbenchmark.rpt—report containing text objects

4 Tables, 17 Fields, 1 Group Field, 10 formatted text objects (text rotation, tool tips, etc.)

245 pages

featuremixbenchmark.rpt

Group #1 Name	
C_NAME	C_ADDRES_PHONE
C_ORDERNO	C_ORDER_DATE
C_CLERK	C_CUST_NAME
C_PRIORITY	C_INTSEGMEN
C_ACCTBAL	C_COMMENT
C_NAME	C_COMMENT
	Sum of CUSTOMER.C_ACCTBAL

summariesbenchmark.rpt

Report Header Section					
Sum of Acct Balance	Sum of CUSTOMER.C_ACCTBAL	Max Acct Balance	Max of	Min Acct Balance	Min of
Sum of Total Price	Sum of ORDER.O_ORDERPRICE	Max Total Price	Max of ORDER.O_ORDERPRICE	Min Total Price	Min of ORDER.O_ORDERPRICE
Total Number Of Orders	Count of ORDER.O_ORDERNO	Distinct Count C Name	CUSTOMER.C_NAME		
Min Order Date	Min of ORDER.O_ORDERDATE	Max Order Date	Max of ORDER.O_ORDERDATE		
Page Header Section					
Group #1 Name	Sum of Acct Balance	Sum of CUSTOMER.C_ACCTBAL	Max Acct Balance	Max of CUSTOMER.C_ACCTBAL	Min of
Group #2 Name	Sum of Total Price	Sum of ORDER.O_ORDERPRICE	Max Total Price	Max of ORDER.O_ORDERPRICE	Min of ORDER.O_ORDERPRICE
Group #3 Name	Total Number Of Orders	Count of CUSTOMER.C_NAME			
Group #4 Name	Total Number Of Customers	Count of ORDER.O_CLERK			
C_NAME	C_ADDRES_PHONE	C_ORDERNO	C_ORDER_DATE	C_CLERK	C_CUST_NAME
	Sum of CUSTOMER.C_ACCTBAL		Distinct Count Order Date		ORDER.O_ORDERDATE
Avg of Acct Balance	Avg of CUSTOMER.C_ACCTBAL		Count Order Date		ORDER.O_ORDERDATE
Distinct of Acct Balance	DistinctCount of CUSTOMER.C_ACCTBAL		Min of Order Date		Min of ORDER.O_ORDERDATE
Avg of Acct Balance	Avg of CUSTOMER.C_ACCTBAL		Max of Order Date		Max of ORDER.O_ORDERDATE
Report Footer Section. All Grand Totals are listed below:					
Sum of Acct Balance	Sum of CUSTOMER.C_ACCTBAL	Max Acct Balance	Max of	Min Acct Balance	Min of
Sum of Total Price	Sum of ORDER.O_ORDERPRICE	Max Total Price	Max of ORDER.O_ORDERPRICE	Min Total Price	Min of ORDER.O_ORDERPRICE
Total Number Of Orders	Count of ORDER.O_ORDERNO	Distinct Count C Name	CUSTOMER.C_NAME		
Min Order Date	Min of ORDER.O_ORDERDATE	Max Order Date	Max of ORDER.O_ORDERDATE		
Page Footer Section					

Appendix 5: Report Data Source

Native IBM DB2 8.1 Server and client using the following TPC-R databases (TPC Transaction Processing Council—www.tpc.org)

Database/Schema Name	Rows	Size KB
tpcsf067	100.000	85.000
tpcsf0067	10.000	7.000
tpcsf00067	1000	

Appendix 6: Software Environment

Crystal Enterprise 10 Configuration and Tests

CE10 Servers	Web App Servers	users	server1 p69001p1	server2 p69001p2	server3 p69001p3	server4 p69001p4	WebSphere
4 x 1CPU	1 x 2CPU	500	2cms, frs, ras	4ps,2cs	4ps,2cs	2js(15)	1*2
4 x 2CPU	1 x 4CPU	1000	4cms, frs, ras	4ps,2cs,1js	4ps,2cs,1js	4ps,2cs,1js	1*2
4 x 4CPU	2 x 4CPU	2000	4cms, frs, ras	8ps,4cs	8ps,4cs	3js(15)	2*3
4 x 8CPU	4 x 4CPU	4000	6cms, frs, ras	16ps,8cs	16ps,8cs	6js(15)	4*2

cms=crystal management server ps=page server js=report server cs=cache server ras=report application server

4 (4 x 1CPU) 1 x 2CPU WAS	8 (4 x 2CPU) 1 x 4CPU WAS	16 (4 x 4CPU) 2 x 4CPU WAS	32 (4 x 8CPU) 4 x 4CPU WAS
			4000 virtual users
			3200 virtual users
		2000 virtual users	2000 virtual users
		1600 virtual users	1600 virtual users
	1000 virtual users	1000 virtual users	1000 virtual users
	800 virtual users	800 virtual users	800 virtual users
600 virtual users	600 virtual users	600 virtual users	
400 virtual users	400 virtual users	400 virtual users	
200 virtual users	200 virtual users		
100 virtual users			

WebSphere

Please refer to the document “Performance Tuning for the Crystal Enterprise Java SDK over the Web” available through the product team for WebSphere performance settings:

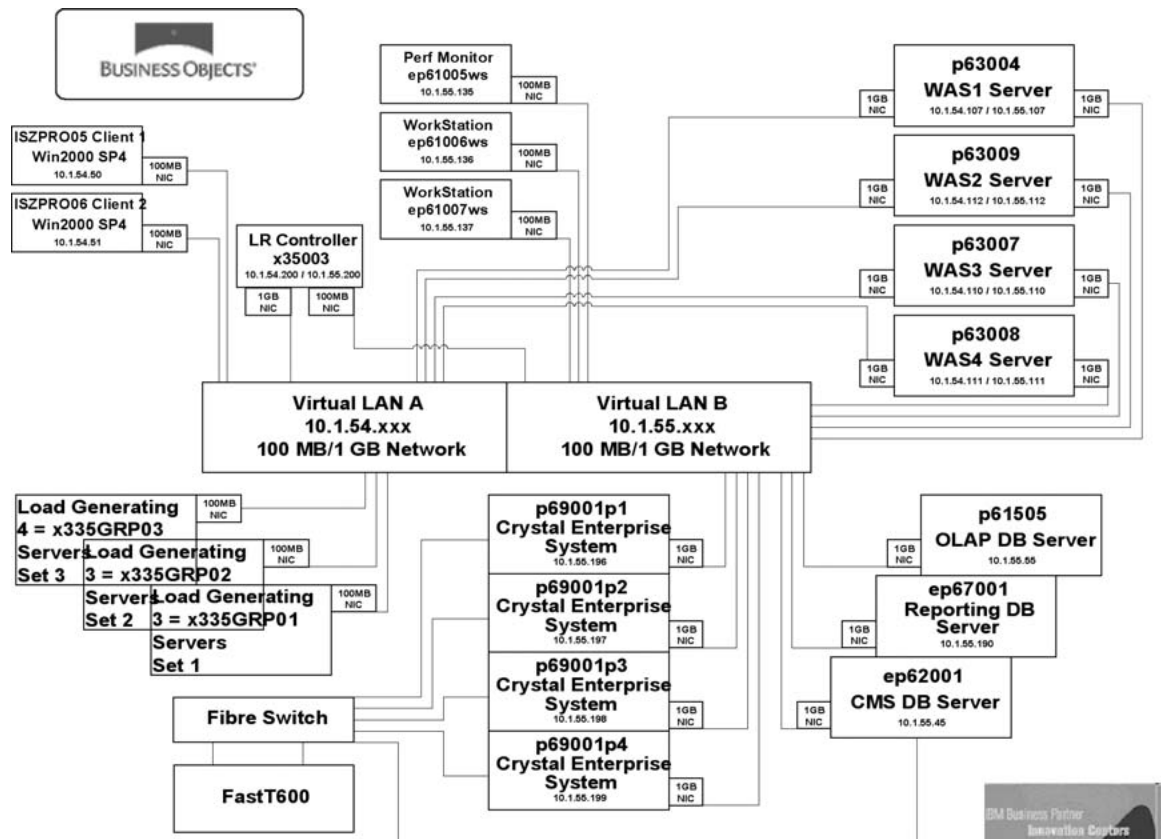
http://salescentral.businessobjects.com/products_services/products/reporting/crystal_enterprise.asp

DB2

Please refer to the document “CE10 System and Database Configuration and Tuning Guide” available through product management for database performance settings:

http://salescentral.businessobjects.com/products_services/products/reporting/crystal_enterprise.asp

Appendix 7: Topology



Appendix 8: Hardware Environment

Crystal Enterprise 10 Servers

- 32 CPU p690
- Gigabit Ethernet
- System Model: IBM, 7040-681
- Processor Type: PowerPC_POWER4
- Number Of Processors: 32
- LPAR Info: 4 LPAR4 Processor Clock Speed: 1300 MHz
- LPAR Info: 4 LPAR4 (p69001p1, p69001p2, p69001p3, p69001p4)
- Memory Size: 32768 MB
- Good Memory Size: 32768 MB

Web Application Servers (CE10 Java SDK and WCA)

3 x 4 CPU p630

Each configured as:

- 10/100 Mbps Ethernet PCI Adapter II
- System Model: IBM, 7028-6E4
- Processor Type: PowerPC_POWER4
- Number Of Processors: 4
- Processor Clock Speed: 1453 MHz
- Memory Size: 8192 MB
- Good Memory Size: 8192 MB

1 x 4 CPU p630

Configured as:

- 10/100 Mbps Ethernet PCI Adapter II
- System Model: IBM, 7028-6C4
- Processor Type: PowerPC_POWER4
- Number Of Processors: 4
- Processor Clock Speed: 1002 MHz
- Memory Size: 4096 MB
- Good Memory Size: 4096 MB

**Americas**

Business Objects Americas
3030 Orchard Parkway
San Jose, California 95134
USA
Tel: +1 408 953 6000
+1 800 877 2340

Asia-Pacific

Business Objects Asia Pacific Pte Ltd
350 Orchard Road
#20-04/06 Shaw House
238868
Singapore
Tel: +65 6887 4228

Europe, Middle East, Africa

Business Objects SA
157-159 rue Anatole France
92309 Levallois-Perret Cedex
France
Tel: +33 1 41 25 21 21

Japan

Business Objects Japan K.K.
Head Office
Yebisu Garden Place Tower 28F
4-20-3 Ebisu, Shibuya-ku
Tokyo 150-6028
Tel: +81 3 5447 3900

For a complete listing of our sales offices, please visit our web site.

► www.businessobjects.com