ACTUATE
The BIRT Company™

Delivering Extreme BIRT iHub Performance

Testing once again proves the superior throughput, scalability and stability of Actuate's embedded analytics and reporting platform

Organizations today place high demands on their embedded analytics and reporting tools. Financial services platforms, mobile environments, and Internet of Things (IoT) deployments are just some of the applications that generate data at high volumes and/or velocities – data that requires analysis, visualization, and reporting to unlock its value. These needs prompted Actuate to push its BIRT iHub product to utilize all available hardware resources.

In a controlled environment, Actuate engineers tested BIRT iHub along several dimensions. They wanted to understand how Actuate's deployment platform, running on bare-metal servers, performed in batch generation and on-demand viewing of reports. The same tests were run on BIRT iHub in clustered environments to see how performance scaled. In separate tests, the engineers ran BIRT iHub at high velocity for an extended period of time on virtual servers, performing almost no maintenance. Results of all of these tests, as documented here, can help companies understand BIRT iHub's strengths and build a technology environment appropriate for their BIRT iHub deployments.

Table of Contents

BIRT iHub 3 was a milestone release for Actuate. Improvements in the product's data model, chart engine, drivers, and many other areas boosted performance, scalability and stability of the software. To accompany the release, Actuate quality assurance engineers thoroughly tested BIRT iHub to validate these improvements and document the software's remarkable capabilities. The tests, based on real-world scenarios, were performed in the first half of 2014 at Actuate's headquarters in San Mateo, CA.

Two sets of tests were run, each with specific goals. **Performance Tests** measured raw system performance of two distinct functions: batch report generation and on-demand viewing. A **Stress Test** placed a heavy, diverse workload on BIRT iHub over an extended period of time. Both tests simulated the operation of a financial services firm that generates and presents statements of transactional data to customers.

This paper details how Actuate designed the tests, and also shares the test results. Companies considering BIRT iHub can use the data in this paper to benchmark BIRT iHub's capabilities against their needs. The test results may also guide hardware selection for companies considering BIRT iHub deployments.

What is BIRT iHub?

BIRT iHub is the commercial platform for all BIRT content. The out-of-the-box deployment server removes the burden on developers to manually build infrastructure components, and automatically enriches content for end users with ineractivity and secure personalization.

Learn more at
actuate.com/products/birt-ihub-visualization-platform/

Test 1: Performance Test

Performance Test Overview

The overall goal of the performance test was to understand the maximum output that BIRT iHub can generate using specific server hardware and clusters of servers. Scenarios were designed to test two separate functions in BIRT iHub:

- **Batch generation** of reports, which is done by the BIRT Factory service.
- **On-demand generation and display** of reports, which is done by the BIRT Online service.

In both sets of scenarios, system performance was measured along two dimensions:

- **Throughput.** How many operations per second could the system perform? How much data per second could the system process?
- **Scaling.** Would performance increase in a predictable, linear fashion as more computing resources were added? Would the system remain responsive as it scaled?

What is the BIRT Factory?

The BIRT Factory is the service in BIRT iHub that handles scheduling of batch reports and generation of BIRT Data Objects. It can be configured to manage resources and optimize performance.

What is BIRT Online?

BIRT Online is the service in BIRT iHub that provides on-demand generation and rendering of reports, interactive crosstabs and dashboards.

Performance Test System Architecture

The architecture assembled for the performance test consisted of:

Servers

Actuate purchased six powerful servers, equivalent to those often used by BIRT iHub customers with demanding applications. Each consisted of:

- A bare-metal server with a 6-core Intel Xeon processor rated at 2.9 Ghz and a SPECint computer benchmark specification rating of 501
- 192 GB of RAM
- 64-bit RedHat Linux Version 6.4
- BIRT iHub 3

Each server was deployed as a single node in a cluster.

What is SPECint?

SPECint is a computer benchmark specification that measures processor power. The benchmark is maintained by the Standard Performance Evaluation Corporation (SPEC). SPECint is the integer performance testing component of the overall SPEC test suite. For more information, see www.spec.org.

Other Components

- Two 600 GB SAS storage systems
- Two Information Consoles to provide portal access to iHub
- Two test clients to generate simulated user actions
- One 10-gigabit network

All test components were assembled as seen in the architecture diagram below.

What is the Information Console?

The Information Console provides an interface that enables users to browse, run, view and organize reports. It seamlessly integrates BIRT content within any webpage, following the site's look and feel.

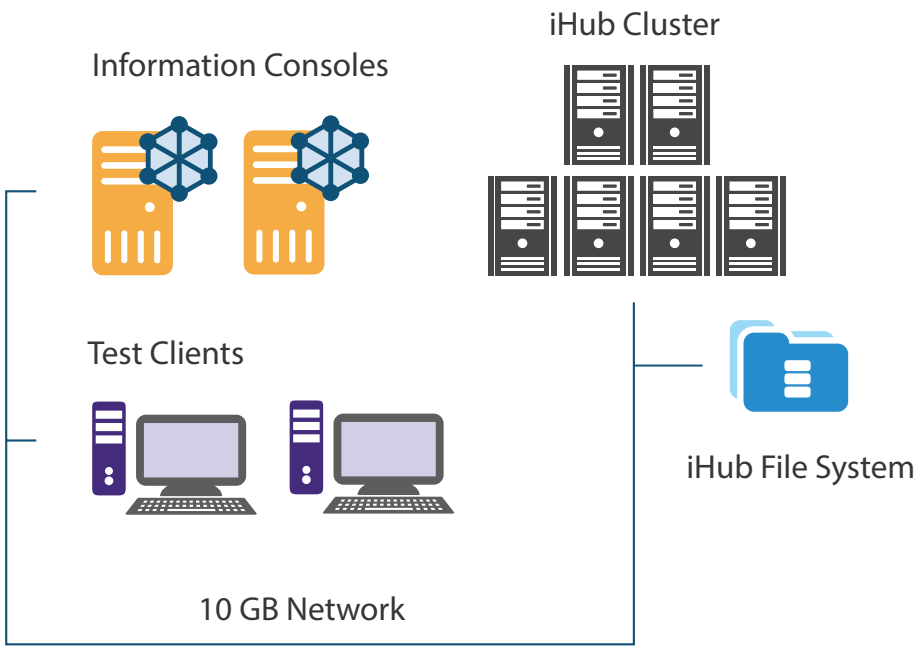

Figure 1: Performance Test Architecture Diagram

Performance Test Scenario 1: Throughput and Scaling of Scheduled Batch Jobs

In BIRT iHub, report designs run in batches of jobs handled by the BIRT Factory service. To test throughput of this service, 50,000 jobs were scheduled to run at one time. Each job created a unique 500-page BIRT document containing a table of transactional data, information about the client, boilerplate text, and a company logo. (The test document is larger than a typical consumer's monthly financial statement.) The BIRT documents were generated for transient online viewing and also stored.

To assess the scaling capabilities of the BIRT Factory service, the test was first run on one server. The same process was then repeated with more servers added as nodes in a BIRT iHub server cluster. No load balancer was used, but test clients were evenly assigned to server nodes.

Performance Test Scenario 1 Results

Nodes	Total time to produce 50,000 reports of 500 pages each	Pages per second produced	MB of reports per second produced	Scaling compared with one node (measured by MB of reports)
1	3 hours 8 minutes	2,215	21.76	N/A
2	1 hour 28 munutes	4,731	46.49	2.13
4	37 minutes 45 seconds	11,027	108.29	4.97
6	29 minutes 35 seconds	14,092	138.39	6.36

A single node of BIRT iHub can produce more than 191 million pages of reports per day (2,215 pages per second times 86,400 seconds in a 24-hour day). A six-node cluster can produce more than 1.2 billion pages of reports. Additional nodes delivered scaling ratios greater than 2:1.

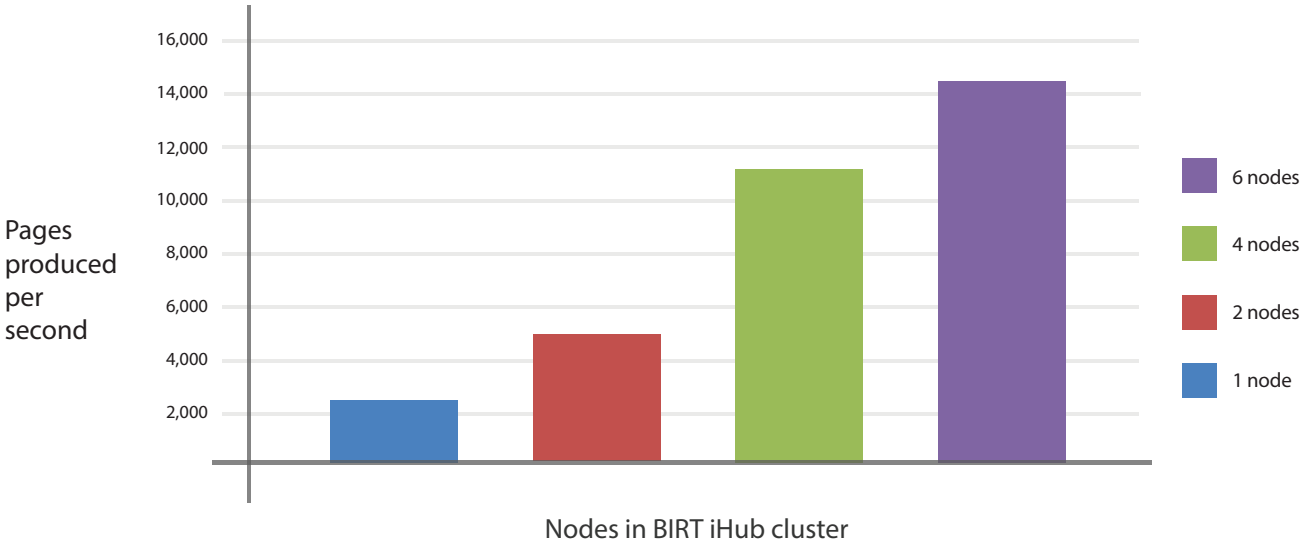

Figure 2: Linear scaling of the BIRT Factory service, which handles batch jobs.

Performance Test Scenario 2: Throughput and Scaling of On-Demand Viewing

In BIRT, on-demand viewing of reports is handled by the BIRT Online service. To test throughput of this service, hundreds of simulated users requested views of 100-page BIRT reports. These reports, generated by BIRT iHub in real time as the users requested them, contained transactional data unique to the user along with standard headers and footers.

To test scaling of the BIRT Online service, the number of users increased and nodes were added to the server cluster at a consistent ratio. No load balancer was used, but test clients were evenly assigned to server nodes.

The single-node configuration was loaded with work from approximately 800 "hot seats" – that is, simultaneous active users cycling in and out of some 800 connections to the BIRT iHub server. Each simulated user generated an action (clicking on a link or interacting with a report) approximately every 30 seconds. The test ran for at least 30 minutes in each of three configurations – one node, two nodes and four nodes.

Performance Test Scenario 2 Results

Nodes	Number of simulated users	Requests fulfilled per second	Scaling of requests fulfilled, compared with one node	Response time to deliver report to user, in seconds
1	800	72.05	N/A	1.329
2	1,600	139.76	1.93	1.365
4	3,200	280.27	3.89	1.382

A single-node BIRT iHub system can handle more than 6.2 million viewing requests per day (72 requests per second times 86,400 seconds in a 24-hour day). Added nodes increased throughput with predictable linear scaling. A requirement of the test – that response time not change significantly as nodes were added and workload increased – was met.

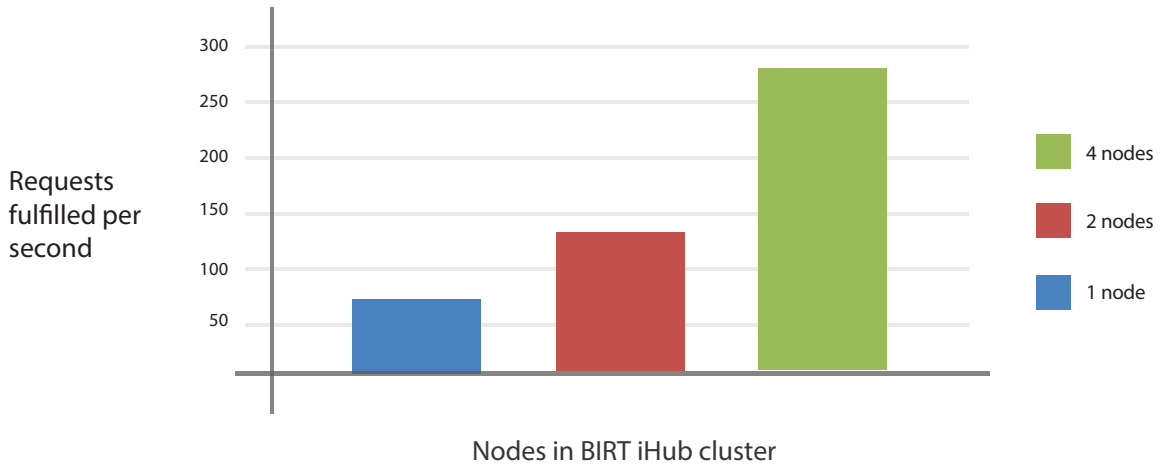

Figure 3: Linear scaling of the BIRT Online service, which handles on-demand viewing.

Test 2: Stress Test

Stress Test Overview

The overall goal of the stress test was to determine how long BIRT iHub could run at a heavy workload without routine maintenance activities. (See "Why Run a Stress Test?") Actuate engineers placed a diverse, high-volume workload on BIRT iHub and maintained that workload over an extended period of time. The stress test uses a large population of simultaneous users, each of whom performs a unique set of operations. Actuate's initial goal was to stress-test its BIRT iHub software for four days – to exceed the equivalent of a long weekend.

Why run a stress test?

Many 24/7 web applications require maintenance activities (such as stopping and restarting compute processes) on a regular basis – sometimes as frequently as weekly. This maintenance frees up system memory and keeps system performance at high levels. In some cases, maintenance windows are used for data backup and archiving or to clear log files. An application that can operate for long periods without maintenance provides administrators with more flexibility in how they manage that application.

Stress Test System Architecture

The architecture assembled for the stress test consisted of:

Servers

The system assembled by Actuate for the stress test consisted of four virtual servers. Virtual servers were used because virtualization is common in enterprise IT shops today, whether servers are on-premise, cloud-based, or a mix of both. Each of the four virtual servers in the test had:

- Two logical CPU cores
- 16 GB RAM
- 64-bit RedHat Linux Version 6.3
- BIRT iHub 3

Other Components

- Two 600 GB SAS storage systems
- Two Information Consoles to provide portal access to the iHub cluster
- Two test clients to generate simulated user actions
- One 10-gigabit network

All test components were arranged according to this architecture diagram.

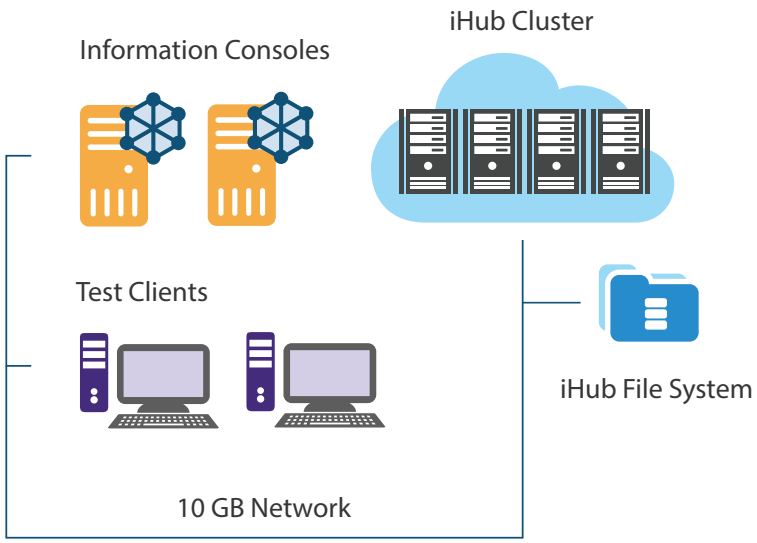

Figure 4: Stress Test Architecture Diagram

Stress Test Scenario

Requests for nine different BIRT iHub operations were generated by simulated users and submitted in rapid succession. These nine operations, both user and administrative tasks, were chosen because they are both resource-intensive and commonly used.

As each operation successfully completed, simulated users requested the same operation again. The rate of operations completed per hour varied according to the operation's complexity. The repeated operations brought CPUs of all four virtual servers to 75 percent of capacity or greater, and held it there for the duration of the test. System response time was not measured in this test.

Operation	Operations per hour	Details of the operation
Run a job immediately	2,300	Jobs included generating and displaying reports
Schedule a new job in the near futrue	3,600	The scheduled jobs were executed within a few minutes
Schedule a new job in the far future	4,800	The scheduled jobs were not executed; the test simply increased the overall job list count
View a report on demand via IDAPI	2,200	Call the Execute Report API directly
View a report on demand via Information Console	34,000	View a BIRT report as a user
View a report document	5,000	View an Actuate BIRT report document file (.rptdocument)
Interact with document	1,100	Perform various Interactive Viewer operations, such as filtering or sorting data in a table
View dashboard	1,400	View various Dashboard operations as defined by user role
Administrative operations	1,680	Select or update users, select roles, and perform other operations using the Access Control List (ACL)

Because the tests ran uninterrupted around the clock, the operational volumes achieved in the stress test were much higher than most, if any, Actuate customers would encounter in a real-world setting.

What is Interactive Viewer?

BIRT Interactive Viewer, included with BIRT iHub, enables everyday users without specialized training to create personalized BIRT content based on their tasks and needs. Built for users who are not familiar or comfortable with traditional reporting tools, BIRT Interactive Viewer allows virtually anyone to access and interact with BIRT content on a self-service basis without affecting the underlying data.

What are BIRT Dashboards?

BIRT Dashboards blend BIRT content (such as charts, crosstabs, tables and BIRT files) with external services in interactive, seamless and configurable web applications. BIRT Dashboards allow self-service authoring and consuming of content by users and enable visual data exploration, analysis, and collaboration.

Extending the Stress Test

After the stress test had run for four days without incident, Actuate engineers considered a different question: How long could BIRT iHub run without interruption or manual intervention? The engineers let the test continue: Simulated users continued to submit requests, and the four virtual servers continued operating at 75 percent or greater CPU capacity. The only manual intervention performed was the occasional moving or deleting of log files, which was done to preserve records in case they were needed for investigation. These actions did not interrupt the BIRT iHub application.

The list of scheduled jobs grew to have millions of entries. This is an important point because the entire list of scheduled jobs must be scanned when BIRT iHub selects jobs for scheduled execution; a longer list requires the more computing resources to scan.

At 30 days, the BIRT iHub system had scheduled 7.5 million jobs, executed 4.2 million jobs, presented 31 million pages of documents for viewing, and performed 1.2 million administrative operations. After more than a month and a half of continuous, uninterrupted operation, Actuate engineers were satisfied with the results and terminated the test. Based on the extreme test workload, engineers estimated that the BIRT iHub test environment had performed the equivalent of a year's worth of activity (based on a more normal workload) during the test period.

Conclusions

Actuate's Performance Tests and Stress Test confirmed that BIRT iHub excels along several dimensions:

Throughput

A dedicated, clustered BIRT iHub system handled workloads at a scale that is rarely seen in the real world. Organizations running high-volume applications – generating millions of documents for millions of users – can expect a well-designed BIRT iHub system to exceed their throughput needs.

Scalability

A dedicated, clustered BIRT iHub system delivered predictable linear scaling as nodes were added to the server cluster. Organizations that expect their needs to grow substantially can add capacity in a predictable manner by increasing the number of nodes in a BIRT iHub cluster.

Stability and Reliability

A modest BIRT iHub system consisting of four virtual servers operated at high volume (greater than 75 percent of CPU capacity) without manual intervention for an extended period of time without crashing, hanging, becoming corrupted, experiencing memory problems or having other issues. Organizations that cannot afford interruptions or downtime can expect unparalleled reliability from BIRT iHub, and can expect response times to remain consistent as a BIRT iHub system scales.

Using the examples and figures in this report, potential customers can plan, select and configure a system architecture to handle workloads of any size. If more users, operations, and/or data are needed, customers can be confident that an iHub-based application can scale via adding nodes to a server cluster, rather than by purchasing costly new hardware.

Global Contacts

Actuate Corporation
Corporate Headquarters
951 Mariner's Island Blvd
San Mateo, CA 94404
Phone: (888) 422-8828
info@actuate.com
www.actuate.com

Actuate UK Ltd.
1st Floor Pickford House
46 Bow Lane
London, EC4M 9DL
United Kingdom
Phone:+44 (0)207 246 4700
ukinfo@actuate.com
www.actuate.com/uk/

Actuate (Deutschland) GmbH
Kaiserstrasse 44
60329 Frankfurt am Main
Germany
Phone:+49(0) 69 66 90 25-0
info@actuate.de
www.actuate.de

Actuate Spain
BIRT Analytics
Frederic Mompou 5
Edificio Euro 3
08960 Sant Just Desvern
Barcelona, Spain
Phone:(+ 34) 93 371 44 70
www.actuate.com/birtanalytics/

Actuate France
57 rue de Châteaudun
Paris 75009
France
Phone:+33 (0) 176 21 5650
fr_contact@actuate.com
www.actuate.fr

Actuate Australia Pty Ltd
Level 10, Suite 7
100 Walker Street
North Sydney
NSW 2060
Australia
Phone: +61 (2) 9900 6300
apac_marketing@actuate.com

Actuate Asia Pacific
8 Temasek Boulevard
#35-01, Suntec Tower Three
038988
Singapore
Phone:+65 6887 3777
apac_marketing@actuate.com
www.actuate.com/ap

Actuate Japan Co., Ltd.
アクチュエイトジャパン株式会社
〒102-0083
東京都千代田区麹町 5 丁目 1
NK真和ビル 9 階
Japan
Phone:(03) 5357-1631
jpinfo@actuate.com
www.actuate.com/jp

Actuate India
Level 8, Vibgyor Towers,
G Block C62,
Bandra Kurla Complex
Mumbai, 400051, India
Phone:+91-22-4090-7281
Fax:+91-22-4090-7283
india_sales@actuate.com

Notice