Experience the power of

Infinite Scalability

Birst's modern web-scale architecture delivers enterprise-grade performance

Delivering analytical capabilities across the enterprise requires supporting large data volumes and high levels of concurrency. This is why Birst has been designed to deliver world-class performance on top of a modern, multi-tenant architecture that can scale infinitely across a shared-nothing, multi-node environment. This means that, unlike traditional client-server based solutions, there is no single bottleneck to hinder performance.

Real-World Usage, Unmatched Performance

During peak production-usage periods, Birst's cloud environment supports **420,000 queries per hour with 77,000 active user concurrency** and an average dashboard response time of 5 seconds. In-house performance lab tests show a total query throughput of over 22,000 queries per hour per node from more than 4,000 active users (20,000 total users) per node with an average response time under 2 seconds. Birst's multi-tenant architecture has been demonstrated to scale linearly (i.e. a four -node configuration will support 4X the workload and users and can process larger data volumes than traditional legacy BI platforms.

Birst Performance and Scalability in Numbers

1.5 Petabytes of data stored*
420,000 queries per hour
10,500,000 dashboards
125,000 dashboard views per day
50,000 data sources
1 multi-tenant architecture

** Only includes data stored in the Birst cloud. Birst customers also analyze massive data volumes stored on-premises.*

A Modern Web-Scale Architecture

Birst is a fully multi-tenant solution that leverages an infinitely scalable pool of computing resources built on a shared-nothing architecture. This stands in stark contrast to traditional client-server based platforms that load data in-memory into a single-server instance which results in performance bottlenecks.

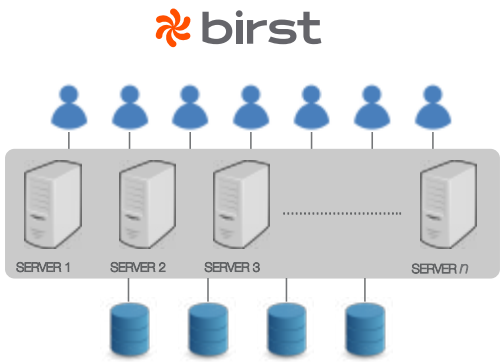

- Modern web-scale architecture
- Multi-tenant environment
- Stateless application servers
- Horizontal scalability

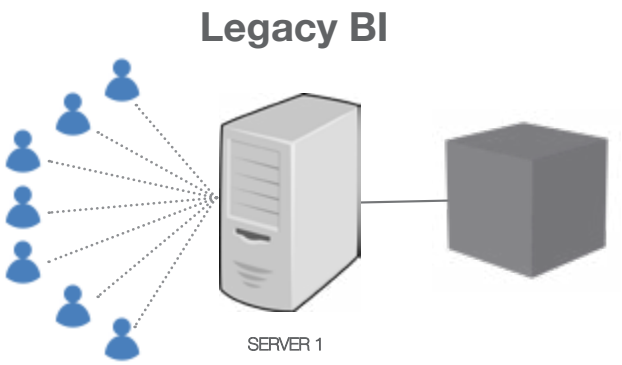

- Client-server clustering
- Single-server dependency
- Proprietary in-memory structures
- Single point of contention (bottleneck)

Birst Linear Scalability Test

Queries/Hour & Active Users with response times < 2 secs

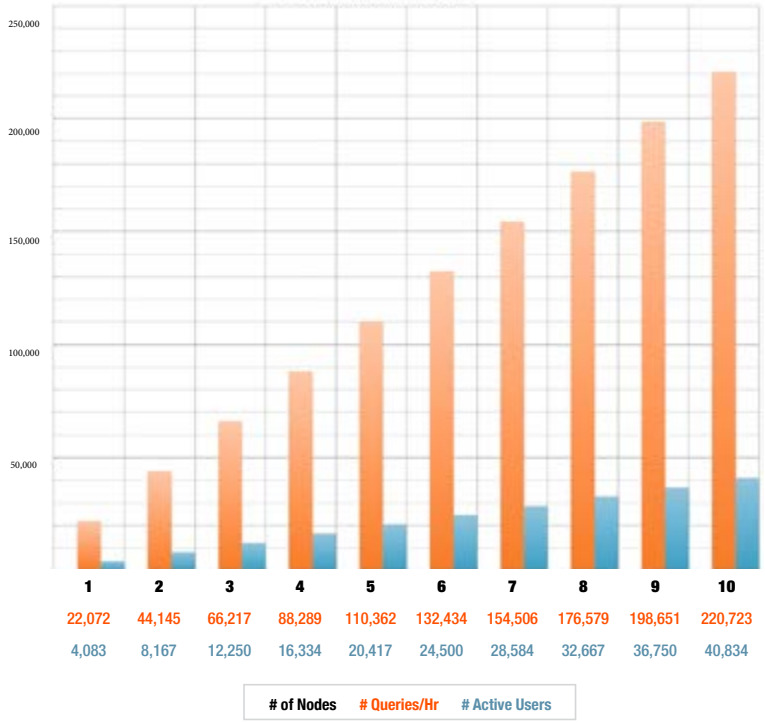

Test Details:

- A series of single and multi-threaded queries were executed against a 4-billion row fact table and multiple dimension tables, the largest containing 0.5 billion records.
- The average response time per query was 1.6 seconds with no degradation up to 10 parallel threads per node.
- On average, each query was comprised of multiple attributes and metrics, or approximately 30 columns.
- The tests assume each active user submits an average of 18.5 queries per hour.
- The tests assume active users represent 20% of total population.
- Testing was performed against a multi-node Amazon Web Services (AWS) environment.
- A single node has 16 virtual CPUs (Intel Xeon® 2.4 GHz) and 64 GB of RAM.

Performance and Scalability Optimizations

- Push-down analytics: Birst ships calculations down to the database platform to leverage its capabilities, while Birst's semantic layer translates operations into database-specific functions.
- Query optimization: Birst generates and optimizes queries – both for data loading and for analysis – appropriate to the backend data source.
- Intelligent multi-tiered caching: Birst provides both in-memory and on-disk caching, using exact and fuzzy cache matching algorithms.
- Globally shared caching: Every query from every user interacts with a common, shared and horizontally scalable caching infrastructure that ensures the best performance in the industry.
- Smart aggregates: Birst can generate and leverage "smart aggregates" to accelerate queries in a more flexible and more scalable way than in-memory cubes.
- "Always On" upgrades and data loads: Thanks to Birst's multi-tenant architecture, users experience no service disruption during system upgrades or data loads. Birst leverages modern load balancers that take individual servers out of the pool and are incrementally upgraded.
- In-database blending: Blended results can be retrieved from any node in the system, eliminating bottlenecks because users are not tied to an individual server.
- Columnar and in-memory storage: Birst offers support for column -store databases (e.g. SQL Server 2014, Amazon Redshift) and in-memory databases (e.g. SAP HANA) for faster and more efficient queries.
- MPP: Birst offers a high performance analytic database option, designed from the ground up for speed. This features massively parallel processing to coordinate processing simultaneously across separate nodes.

Reckitt Benckiser · Kellogg's · CITRIX · American Express · CROSSMARK · Principal Financial Group · at&t

birst

Call toll free: (866) 940-1496
Email us: info@birst.com
www.birst.com

About Birst

Birst is the global leader in Cloud BI and Analytics. The company helps organizations make thousands of decisions better, every day, for every person. Birst's patent-pending 2-tier data analytics and BI platform enables enterprises to create a trusted source of data, place it in the context of key business users and then enable business users up and down the organization to report and analyze the information using world-class BI tools. Thousands of the most demanding businesses trust Birst to make metric-driven business execution a reality. Learn more at www.birst.com and join the conversation @birstbi.