



A modular, affordable approach for better business decisions

Sun Business Intelligence Solution for Web Infrastructure

Your customer data is one of your greatest assets, especially on the Web. But to get maximal value from that information, you need to be able to make connections between back-office data and ongoing Web-user behavior in near real time. Only then can you create the kinds of deep customer relationships and rich online experiences that keep users coming back.

But until now, those kinds of analytics solutions — a combination of data warehousing and BI solutions — were proprietary and expensive. And that

that make customers feel welcome and remembered, and give your sales unit the knowledge it needs to engage with customers more effectively. And the same solution will also give you a broad understanding of your business, so you can make better decisions for your company.

Infobright: innovative data warehousing

Infobright Enterprise Edition (IEE) is a commercial open-source data warehousing software solution that delivers high

This unified solution can dramatically lower the cost of entry to create a powerful, sophisticated analytics solution.

Highlights

- Reduce the cost and complexity of deploying a sophisticated business intelligence (BI) solution
- Combine Sun Open Storage with commercial open-source solutions: Pentaho, for business intelligence; and Infobright, built with MySQL™ software, for data warehousing
- Choose between two recommended solution levels to fit your data requirements and your budget
- Save up to 75% versus the cost of proprietary storage, and up to 90% off the cost of proprietary Web analytics software

kept all but the largest enterprises from moving forward.

Sun Business Intelligence Solution for Web Infrastructure is changing all that by combining Sun's Open Network Systems with commercial open-source solutions: Pentaho, for the BI and extraction, transformation, and loading (ETL) functions; and Infobright, for data warehousing. This unified solution can dramatically lower the cost of entry to create a powerful, sophisticated analytics solution.

So now you can do what all the largest and most successful sites do — create the kinds of personalized online experiences

performance, reduced administration requirements, and low TCO. Because it combines a compressed, column-oriented database with a unique "Knowledge Grid" architecture, it can deliver near-real-time responses to complex queries using dynamic data, so it's perfect for Web analytics deployments.

At its core is the world's most popular open-source relational database management system, MySQL software — a fully integrated, transaction-safe, ACID-compliant database — with full commit, rollback, crash-recovery, and row-level locking capabilities, plus master/slave replication configurations to increase read scalability across larger server farms.

And unlike traditional data warehousing solutions that require extensive IT support to construct and deploy databases, Infobright is self managing, with no need to create indexes or materialized views, or partition data, so you can reduce administration costs by up to 90%.

Its average 40:1 data compression lowers your storage capacity requirements. And it provides the features you need to manage your data effectively:

- Rapid data load times — up to 280 GB/hour
- Constant query and load performance, even as the database grows
- Scalability up to 30 TB or more of compressed data

Pentaho: proven BI

As the world's most popular open-source BI software, Pentaho provides a complete suite of tools for enterprise reporting, analysis, dashboard, data mining, workflow, and data integration. Its products are being used by many leading organizations, including the U.S. Naval Air Systems Command, Life-time Networks, and Terra Industries.

Its user support was rated best in the industry by a top analyst, and it's the only open-source BI suite provider to lead and sponsor all of its projects, a critical factor in delivering the service-level agreements (SLAs) that customers demand.

You can use the complete suite, or pick pieces to match your specific business requirements:

- Pentaho Reporting enables you to easily access, format, and distribute information to employees, customers, and partners.

- Pentaho Analysis helps you gain the insights and understanding you need to make optimal business decisions.
- Pentaho Dashboards deliver key metrics on individual, departmental, or enterprise performance in an intuitive visual interface, so you can understand and improve your organization.
- Pentaho Data Integration delivers powerful ETL capabilities using an innovative, metadata-driven approach.
- Pentaho Data Mining helps you uncover meaningful patterns and correlations that might otherwise remain hidden, by running data through sophisticated algorithms.

Increased performance and value through integration

While each of these software solutions is effective on its own, by putting them together and deploying them on Sun Open Storage you can create a powerful Web analytics solution for significantly less than you'd pay for a proprietary solution. The Infobright and Pentaho software components each can save you up to 90% versus the cost of proprietary solutions, while Sun Open Storage provides savings of up to 75% versus traditional proprietary storage solutions, while consuming up to 3x less power.

Since Infobright is built upon MySQL software, you can reduce maintenance requirements because the Infobright software leverages extensive tool connectivity provided by MySQL connectors, including C, Java Database Connectivity™ (JDBC), Open Database Connectivity (ODBC), .NET, and Perl. And running Infobright software on a Sun™ Storage Unified Storage System demonstrated

up to a 92% improvement in query performance compared to a competitive storage device.

Fit the solution to your needs

To simplify and speed deployment, Sun Business Intelligence Solution for Web Infrastructure has two recommended configurations, with different hardware components to suit varying data requirements and budgets.

The entry-level solution combines Pentaho and Infobright software with the Sun Fire™ X4275 server, which combines both storage and compute capabilities and delivers the best performance in its class for data warehousing applications. With twelve 3.5 in. drive bays that can support SAS or SATA hard disk drives (HDDs), or solid-state drive (SSD) disks, it can hold up to 12 TB of data internally, and will accommodate up to 24 TB in the near future. Additional storage can be added by utilizing the Sun Storage J4000 array series of disk enclosures.

What's more, it's flexible because of its open architecture, so you can deploy with any operating system (OS), including Windows™, Linux, and the Solaris™ OS. With the new Intel® Xeon® Processor 5500 series, you can get up to 3x the performance available with the previous generation of Intel Processors. And because Intel and Sun collaborated during the engineering and design of the new processors, you'll get additional performance, efficiency, and reliability benefits if you deploy with the Solaris OS. You can boost performance further by deploying new Flash-based SSD drives, achieving up to 65x better response times and 7x better throughput over hard disks.

It all adds up to exceptional compute performance and storage that costs less than \$1/GB, making it the perfect backbone for a high-performance Web analytics solution for enterprises with limited data requirements. And with the modular approach inherent in the system, you can add additional storage and compute power as your needs change.

Enterprise-class storage, big-time performance

For organizations with higher data requirements, the same combination of Pentaho and Infobright software can be deployed on a Sun Storage 7310 Unified Storage System. This upper mid-level device in the Sun Storage 7000 Unified Storage System series enables you to scale from 12 TB up to 30 TB of data, depending on your needs, while saving you 75% versus proprietary storage solutions. The device itself scales up to 96 TB.

It integrates a number of powerful features for breakthrough speed and scalability, simplified administration, and unprecedented storage analytics. First, it employs Hybrid Storage Pool technology, unique to Sun systems, which uses Solaris ZFS™ File System to integrate Flash-based SSDs, with much faster access times and reduced power needs, in both read and write caches. It also intelligently decides what information to stage in the appropriate cache, depending on the functions you're performing. It's a seamless, integrated approach that can boost I/O performance up to 3.2x, with no intervention required by an administrator, and boost capacity up to 2x, while reducing floor space and using up to 4x less power.

Solaris ZFS File System, built into the Sun Storage 7310 Unified Storage System, reduces the time required to perform many functions, from hours to just minutes, and enables you to automate administrative tasks — saving time and money. What's more, it automatically performs error checking, and repairs corrupted data, to deliver the most comprehensive self-healing capabilities available.

The Sun Storage 7310 Unified Storage System also boasts unprecedented storage analytics capabilities that give you real-time visibility through the entire data path, so you can analyze performance issues down to the user level; see the amount of memory, CPU time, filesystem, and network resources used by the active processes; and look at real-time statistical graphs to help locate problems. And with preemptive issue resolution, it offers pointers on how to increase performance.

It is also simple to install and manage. It can be configured in less than five minutes, in six simple steps. A pre-installed installation wizard takes the guesswork out of tuning, and preloaded storage configuration profiles match the hardware configuration to your workload. The system's GUI makes managing your storage intuitive — you can see your system easily and clearly, and modify things with click-and-drag administration.

The Sun Storage 7310 Unified Storage System can cluster to deliver high availability for mission-critical implementations. And if you need to scale beyond its storage capacity, the Sun Storage 7410 Unified Storage System provides the same functionality, plus up to 288 TB of storage, soon to be up to 576 TB.

Boosting compute power for high data requirements

Depending on your data requirements, user traffic, and workload, you might require separate compute and storage nodes, and might even want to separate the data warehousing and BI functions onto separate compute nodes. Sun provides this flexibility with several servers that can meet these requirements at the right price.

The Sun Fire™ X2270 server is the industry's first Flash-ready server from a Tier 1 vendor. It delivers the best price/performance, top energy savings, and the most capacity of any comparable system in its class. It can be scaled from four to eight processing cores in a 1U form factor, and has four drive bays for up to 4 TB of internal storage.

The Sun Fire™ X4170 server is a versatile enterprise server that delivers the performance of many 4U systems in a compact 1U form factor — saving on space, power, and cooling requirements. And it costs half the price of competing systems.

Most importantly, because the architecture is flexible and the components are interoperable, you always have the freedom to scale how and when you need to — with the servers and storage devices that fit your business needs.

The services to support your deployment

Sun Professional ServicesSM also provides a comprehensive set of services to help you get the most from your analytics solution, mitigating the risk of downtime, delay, and data loss.

- Sun Implementation and Assessment Services help you go beyond standard setup and the basic diagnostics of standard installation, providing select configuration and validation to ensure that your systems are running at an optimal level right from the start.
- SunSM Unified Storage Data Migration Services help you minimize disruption and achieve full functionality of new technology investments. Sun's deep technical expertise and proven implementation and migration tools help you speed ROI by securely migrating your data, utilizing MD5 hashing for file and security validation.
- Sun Single Point of Contact (SPOC) Service provides comprehensive management of IT needs on an ongoing basis, including case creation for all incidents, dispatches to the appropriate support vendor, and case monitoring according to the responsible vendor's contract terms. So you can focus on what you do best, and let Sun handle any technical problems.
- SunSpectrumSM Service Plan simplifies and reduces the cost of managing an IT infrastructure while delivering proven ROI.

- Sun Learning Services help administrators gain a deep knowledge of Sun Open Network systems and enable staff to be more productive.

The intelligence to act

You know the kind of Web experience you want to create for your users — one that is responsive to their needs and interests, that remembers what they've done and how they've used your service before, and that works with them, offering them the choices and the services that are most appropriate. To create this, you need effective analytics that combine powerful BI and data warehousing solutions, while keeping costs in line.

The combination of Sun Open Storage and the commercial open-source Infobright data warehousing and Pentaho business intelligence and ETL software packages dramatically lowers the cost of entry and simplifies deployment of a powerful, sophisticated BI infrastructure. So you can connect data from business functions with Web viewing behavior to create the kind of user experience that energizes sales and boosts customer loyalty — without breaking the bank.

To learn more, go to:
sun.com/WebBI