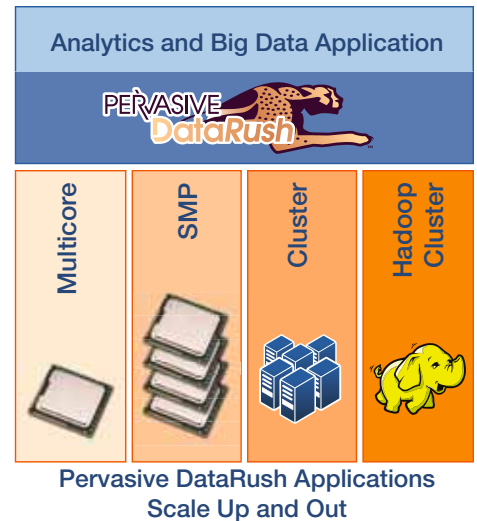


Scale Your Analytics and Big Data Application Across Many Cores and Multiple Nodes



Impact of Pervasive DataRush™...

- Deliver high-throughput analytics on massive datasets
- Accelerate every node in your Apache Hadoop cluster
- Eliminate performance limitations of production Python apps
- Scale your big data application across many cores and multiple nodes
- Build apps that scale from 2 to 384 cores... and counting



Industry Trends

Big Data

If data is a critical asset in your organization, you have likely experienced the trend to bigger and more complex data that is transforming how you process and analyze data. The impact can be significant:

- performance bottlenecks have emerged
- you've been forced to sample data to reduce analysis time
- you've had to reset performance expectations
- you're considering a complex development project to utilize a large cluster.

Fortunately, there is a better way.

Multicore Servers

The dramatic shift by Intel, AMD, IBM and Sun to multicore processors provides a compelling approach to dealing with big data. This means that you now have a large and growing number of cores available to you on a single server. So, if you are able to leverage the power of every core and scale over time as additional cores are added, you get a low-cost, high-performance approach to solving the big data problem.

And by accomplishing more on a single server or each node in your cluster, you reduce the footprint and power requirements, providing a greener approach to handling big data.

Why Pervasive DataRush?

- **Scalable.** Dynamically scale your performance with increased core counts or increased nodes. Add more cores or more nodes and your application performance increases!
- **High-Throughput.** Parallel dataflow technology enables fast, deep analysis of large data sets with no limit to input size.
- **Cost Efficient.** Get the highest possible performance from commodity multicore servers and SMP systems.
- **Easy to Implement.** Allows fast implementation while hiding the complexity of parallel programming.
- **Extensible.** A fully extensible platform so you remain in control of development.

Pervasive DataRush is a parallel dataflow platform that eliminates bottlenecks in your data-intensive applications

Pervasive DataRush for Data Preparation

- Improve data standardization
- Increase performance when relational databases are used for data processing
- Leverage fuzzy matching of fields
- Improve de-duplication of records
- Audit all incoming data
- Find linkages between transactions and the master database
- Perform retrospective analysis of large data repositories
- Reduce cost and complexity of deployment by fully leveraging multicore servers

Pervasive DataRush for Analytics

- Improve turnaround times on analysis
- Reduce cost and performance impact of in-database analytics
- Discover more anomalous patterns (fraud, intrusions)
- Refresh models more often
- Stop sampling and use all data
- Predict outcomes such as churn and enrollment
- Reduce time spent preparing data
- Reduce cost and complexity of deployment by fully leveraging multicore servers

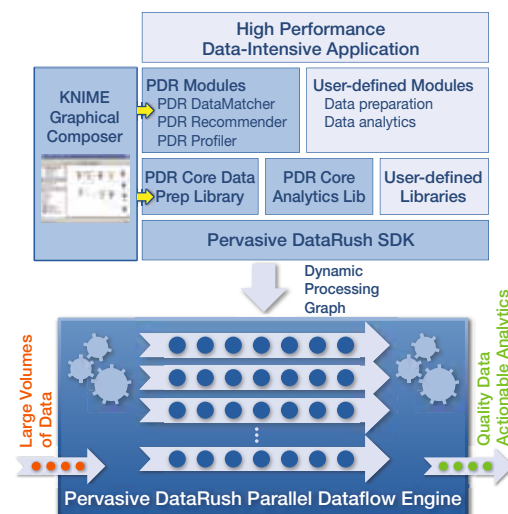
Pervasive DataRush Products

- Pervasive DataRush Parallel Dataflow Platform
- Pervasive DataRush SDK
- Pervasive DataRush Core Libraries
 - Data Preparation
 - Analytics
- Pervasive DataRush Profiler
- Pervasive DataRush for KNIME
- Pervasive DataRush DataMatcher
- Pervasive DataRush Recommender

Pervasive DataRush Parallel Dataflow Engine

At the heart of every Pervasive DataRush-enabled application is the Pervasive DataRush Parallel Dataflow Engine. This embedded engine transforms the way an application handles large volumes of data. The Pervasive DataRush engine is designed to take full advantage of the disruptive shift to multicore processors, enabling performance on a single server or small cluster that you thought was only possible on a very large, complicated cluster. Designed to handle gigabytes to petabytes of data, the Pervasive DataRush engine scales with both the size and the dimensionality of the data.

Pervasive DataRush supports a wide array of operating systems including Windows, Solaris, IBM AIX, Linux and HP-UX. In addition, Pervasive DataRush supports multicore CPUs from AMD, IBM, Intel and Sun.



About Pervasive Software

For over two decades, Pervasive has helped thousands of customers around the world in every industry to manage, integrate and analyze their data. Pervasive DataRush is a parallel dataflow platform enabling data-intensive applications such as claims processing, risk analysis, fraud detection, data mining, predictive analytics, bioinformatics and life sciences.

Contact Us

www.persivedatarush.com
1-866-980-RUSH (7874)
1-512-231-6000 (Main Office)

info@persivedatarush.com