

A White Paper

Real-World Strategies for Big Data

**Tackling The Most Common Challenges
With Big Data Integration**



Information Builders helps organizations transform data into business value. Our business intelligence, integration, and data integrity solutions enable smarter decision-making, strengthen customer relationships, improve performance, and drive growth.

Table of Contents

1	Introduction
2	How to Make Your Big Data Strategy Succeed
3	Technical Challenges in Big Data Implementations
6	Uniting Strategy and Technology
6	Key Capabilities
9	Conclusion
10	Glossary

Introduction

More organizations today are realizing that there is tremendous value locked away within their big data, and are seeking to tap into it to improve operations, boost revenues, and gain a competitive edge. When asked to list their top IT priorities in a recent survey, only 6 percent of companies didn't include big data in their top ten.¹

Companies that approach their big data strategically reap countless benefits. They have a deeper understanding of their customers, make better and more informed decisions, and are more agile and competitive. They also generate 69 percent more revenue per employee, according to the consulting firm Perficient.

Perhaps that's why IDC expects big data technology and services to grow at a 23.1 percent compound annual growth rate, reaching \$48.6 billion in 2019.²

Due to the many technical obstacles that can arise, a majority of big data projects fail. The complexity of Hadoop®, a shortage of skilled talent, the potential for cost overruns, and tools that lack scalability, security, and the ability to manage data quality are among the many reasons why big data initiatives don't deliver the desired results, or stall completely without ever being completed.

This white paper will explore the technical challenges that hinder big data success. It will also explain what you need to succeed and highlight iWay Big Data Integrator, a modern, native solution for Hadoop-based data integration and management that ensures high levels of capability, compatibility, and flexibility.



¹ Narisi, Sam. "Why 55% of Big Data Projects Fail," IT Manager Daily, February 2013.

² "Six Predictions for Big Data Analytics and Cognitive Computing in 2016," Bloomberg, January 2016.

How to Make Your Big Data Strategy Succeed

Because big data platforms are often seen as the solution to everyone's information challenges, organizations have a tendency to apply Hadoop as a replacement for all traditional systems. But while Hadoop offers many advanced capabilities, applying it to use cases for which it isn't suited will amplify its complexity and derail your big data project. Start with use cases that can't be addressed with a traditional data warehouse, or will be less costly when done in Hadoop.

The most successful big data strategies include four key elements:

People

Big data strategies require a mix of skill sets from across the organization to ensure that the appropriate knowledge exists for each facet of the project. You should include your data scientists, as well as data warehousing teams, data quality specialists, and other people with heavy tech expertise.

Intent

Make sure you understand what you are looking to accomplish at the outset. Do you want to create new processes for information management? Drive better decision-making? Make it easier for users to interact with large data sets that they are unable to access in other systems? Capture streaming content and turn it into usable information? Working with your goals always in mind will keep your strategy on the right track.

Data

What data assets should you deploy to Hadoop? That depends on your end users. It is important to arrange your big data so it is easily consumable and usable to those who need it most. This, however, can be time-consuming. Studies claim that data scientists spend as much as 80 percent of their time preparing and managing data for analysis. This includes collecting, cleaning, and organizing data; and 76 percent of these data scientists view these tasks as the least enjoyable part of their work.³ A smart approach to data ingestion and preparation can save a lot of time and many headaches.

Tools

The right supporting tools are the key to successfully executing your big data strategy. Many solutions provide functionality that improves Hadoop utilization. However, these tools vary in complexity and serve to accomplish different goals or address different specific problems. For example, some are command-line driven and some are Linux-based, requiring special expertise. Others offer a solution to some big data problems, but require leveraging additional tools to solve others. Most call for specific knowledge and/or extensive training. Choose wisely.

Information Builders' iWay Big Data Integrator (BDI) provides a simple interface that marshals the most common Hadoop resources. Most importantly, data does not need to be copied or moved outside of Hadoop to be cleansed with BDI, and there is not a data maintenance burden across two systems. BDI can natively ingest, transform, cleanse, and manage data lakes. It supplies a modern, native approach to Hadoop-based data integration and management.

³ Press, Gil. "Cleaning Big Data: Most Time-Consuming, Least Enjoyable Data Science Task, Survey Says," Forbes, March 2016.

Technical Challenges in Big Data Implementations

Despite their importance, more than half (55 percent) of big data projects are never completed.⁴ While the reasons for failure vary from project to project and company to company, there are some common technical obstacles to be on the lookout for. Preparing for and/or avoiding these challenges can dramatically increase the chances of your initiative's success.

Hadoop Complexity

While Hadoop deserves the praise it gets for its ability to handle massive amounts of data, it is notoriously difficult to work with. In a recent poll, 73 percent of Hadoop users said that understanding the platform was the biggest challenge of their project.⁵ Hadoop is not for your average programmer – it requires expertise in Java™, as well as a deep understanding of Python, Scala, R, and more. As a result, your company may find itself focused on getting familiar with the technology, rather than critical big data use cases.

iWay BDI removes the complexities of ingesting, transforming, and removing data simplifying the creation, management, and use of Hadoop-based data lakes. With its straightforward interface, users don't need to concern themselves with Hadoop technologies and intricacies. Users can focus on working with the data, rather than learning the technology.

Shortage of Talent

Big data projects call for a sophisticated team made up of developers, data scientists, analytics experts, and data governance professionals. Surveys show that 80 percent of CIOs struggle to find experienced and knowledgeable Hadoop talent to tackle the numerous tasks and steps involved in Hadoop applications.⁶ Many big data solution vendors try to overcome this shortage by selling educational services, or by creating products to handle specific problems such as data ingestion or preparation. But these fall short because they don't address end-to-end Hadoop utilization.

Hadoop is complex and requires specialized skillsets to deploy simple data flows. With BDI, users who are trained in data warehousing and ETL can quickly apply these same skillsets to Hadoop. By applying a familiar data warehousing paradigm on top of Hadoop, users can create data processing pipelines without having to be re-trained on big data technologies.

Lack of Scalability

You're dealing with enormous amount of data, with volumes growing by the day. The scope of your project, and the hardware and resources needed, will evolve throughout the course of your initiative, and can quickly get out of control. Furthermore, resources are often needed in "bursts." There will be short periods of time when you're consuming all your available resources, and brief periods when you need almost none. This makes it difficult to allocate processing power and resources for the necessary jobs.

iWay Big Data Integrator has high scalability and performance. It can support all your data whether it is traditional data, machine-generated data, social media, web data, or data from the Internet of Things (IoT).

⁴ Narisi, Sam. "Why 55% of Big Data Project Fail," IT Manager Daily, February 2013.

⁵ Azunga, "The Challenges of Big Data and How to Work Around Them," Flagship Solutions Group, May 2016.

⁶ "CIOs and Big Data," Visually, Inc. January 2013.



Actionable Insights

Once the key goals of your big data plan have been identified, you need to make sure you can ingest the needed data to glean the desired insights in a way that can be acted upon. What sources do you need to collect? Are there other sources you can incorporate to further extend value? How do these sources relate to each other? How will you join them together?

Creating actionable business insight requires data from many disparate data sources. Sometimes, specialized technical teams are needed to add new data sources to data lakes, and transformation of that data into a usable form requires additional effort. With iWay BDI, adding the new data source is simple and transformations can be visually edited for inclusion into target outputs.

Poor Data Quality

Ovum Research claims that companies lose an estimated 30 percent of their revenues due to poor data quality.⁷ The potential for quality issues in big data scenarios is quite high, given the high volume of data being processed, the vast number of sources it comes from, and the varied formats it exists in. A recent big data survey reveals that 60 percent of IT leaders believe their organizations lack accountability for data quality, while half question the validity of their data.⁸

However, cleansing and standardizing massive data sets can be an overwhelming task – especially if data quality is approached manually. The right supporting technologies can automate profiling, cleansing, and other critical data quality management processes to optimize accuracy and consistency at all times.

Most data sources have data quality problems, even if it's your master data. Changing this source data is costly and time-consuming process. With BDI, sources can be consumed as they are, and data quality rules can be used for data profiling, analysis, enrichment, merging, and cleansing from native Hadoop sources and file formats. This standardizing of data provides a practical format for both operational and analytical usage.

Lack of Data Security

Hadoop clusters have no security protocols in place by default. Yet, certain information assets are sensitive or confidential in nature, and need to be protected from unauthorized use. You'll likely need to use encryption for both data at rest and data in motion, especially if it is coming from or stored in the cloud.

Hadoop alone does not enforce security, and as a result, technologies such as Apache Sentry, Knox, and Ranger have been developed to impose role-based access to Hadoop resources. iWay Big Data Integrator obeys Hadoop cluster policies as defined by Apache Sentry, Knox, and Ranger. Your organization does not need to be concerned about unauthorized team members accessing, modifying, or deleting information because locations and repositories are secured.

⁷ Lawson, Loraine. "What Does Bad Data Cost," IT Business Edge, February 2015.

⁸ Shacklett, Mary. Data Quality: "The Ugly Duckling of Big Data," TechRepublic, February 2014.

Cost Management

With expensive and hard-to-find talent, complex tools, specialized resources, data quality problems, and other issues, the cost of your big data project can add up quite quickly. Keep potential costs in mind as you create your big data plan.

Big data deployment costs are difficult to manage without the proper toolset. iWay BDI helps manage the cost of development by minimizing the need for specialized talent and technical expertise. Implementing a data quality program often causes re-work, but can be controlled by applying data quality rules to the data that's already ingested. Security breaches to sensitive data are minimized by applying the appropriate security rules. With this toolset, Hadoop users and groups can focus on solving use cases instead of technical problems.



Uniting Strategy and Technology

With iWay Big Data Integrator, Information Builders helps customers minimize the complexities associated with big data initiatives, providing a modern, native approach to Hadoop-based data integration and management.

Benefits of iWay Big Data Integrator:

- Provides a simplified, easy-to-use interface to perform data integration in Hadoop, including data ingestion, transformation, and data quality. Users won't need as much training or Hadoop knowledge to be effective
- Eliminates Hadoop coding, enabling a more rapid response to business requirements using less-expensive skill sets
- Runs under YARN, taking advantage of native Hadoop performance and resource negotiation
- Marshals Hadoop resources, e.g., Sqoop and Flume, and leverages Hadoop-standard file formats, such as Avro, to support any kind of big data integration use case
- Ingests non-Hadoop data (transactions, messages, alternate data sources) through iWay Service Manager
- Includes additional sophisticated process management and governance
- Enables integration processes to run as MapReduce or Spark applications or a combination of these with a single configuration parameter
- Integrates industry-accepted components of the Hadoop ecosystem, ensuring long-term stability in dynamic environments
- Runs on all major Hadoop distributions
- Transforms big data quality by leveraging Hadoop implementations to improve operational excellence

Key Capabilities

Data Ingestion

iWay Big Data Integrator hides the complexity of data ingestion, replacement, and de-duplication using Sqoop, Flume, Kafka, Spark, and our iWay Service Manager – without programming.

- Easily ingest any JDBC-accessible data into Hadoop Distributed File System (HDFS) using Sqoop
 - Source data and metadata is ingested and exposed in HCatalog for efficient data processing in Hadoop.
 - Changed Data Capture (CDC) is added since Hadoop currently doesn't offer a native CDC capability
 - Mainframe dataset records are ingested into HDFS, using Sqoop as the underlying technology
- Leverage Apache Flume, Kafka, Spark, and/or iWay Service Manager for unstructured and

streaming data

- Transport streaming data sources (e.g., network traffic data, social media data, log data, etc.) into HDFS using Apache Flume as the underlying technology. It allows for Flume script validation and assigns metadata using a data wrangling user interface
- Out-of-the-box morphlines enable rapid deployment of JSON- and XML-to-Avro transformations
- Customizable Spark stream ingestion, configurable through Scala and Java

In addition, iWay Service Manager can run natively within the cluster, allowing the Hadoop administrator to manage the ingestion of multiple sources of legacy data (e.g., transactions, messages, data sources, etc.). iWay Service Manager also provides native data transformation capabilities with a rich array of out-of-the-box functions.

Data Transformation

After ingestion, iWay BDI can transform unaltered data into integrated data sets to support downstream business processes. It leverages the Hadoop ecosystem and native components to transform data from many sources to one or many targets, which can be sources for follow-on transformations. BDI exposes all operators and functions that are part of the individual Hadoop distribution.

iWay Big Data Integrator also provides an SQL analysis and migration feature that lets users transform all data that resides on Hadoop. This allows end users to leverage Spark or MapReduce to integrate data and create new content.

BDI contains a complete Spark development environment that manages Spark dependencies, build process, and publishing to the chosen cluster node. For transformations that are too complex for standard SQL representation, users can leverage Scala or Java to manipulate data.

All newly integrated content remains natively in Hadoop and can be published to traditional data repositories.

Predictive Modeling/Data Science

When Apache Spark is fully integrated into iWay Big Data Integrator, all of Spark becomes available. MLlib, Spark's scalable machine learning library, has capabilities that can be used by Data Scientists and Analysts. They can create quantitative and qualitative analyses all from within the product. Dependencies, build processes, and publishing to cluster nodes are also handled from within BDI.

iWay Big Data Integrator can leverage Spark GraphX to create content for graph database and relationship analyses.



Collaboration

iWay Big Data Integrator helps to enable multiple facets of the organization:

- Data warehousing personnel can create data ingestion and publishing processes with no knowledge of the underlying technology
- ETL professionals can transform content on Hadoop through Spark or MapReduce without significant training
- Data scientists can consume this content to produce actionable insights

These allow the organization to cross-functionally design and deploy big data processes in an agile fashion.

iWay Big Data Integrator Editions

Two editions make the right components available to organizations at the right price. Both editions run under YARN and leverage YARN-managed resources and MapReduce technology.

- The Integration Edition provides JDBC-accessible data replication and changed data capture, data de-duplication, and streaming and unstructured data capture
- The Data Quality Edition (runs under Spark only) includes everything in the Integration Edition, and adds data profiling, analysis, enrichment, merging, and cleansing from native Hadoop sources and file formats (Hive, Avro, sequence file, Snappy-compressed CSV, etc.)



Conclusion



“With traditional solutions becoming too expensive to scale or adapt to rapidly evolving conditions, companies are scrambling to find affordable technologies that will help them store, process, and query all of their data,” analyst firm Forrester says. “Innovative solutions will enable companies to extract maximum value from big data and create differentiated, more personal customer experiences.”⁹

In the past, technical obstacles such as lack of scalability, ineffective security, and shortage of talent have prevented organizations from truly benefitting from their big data. Now, with Information Builders, big data integration and management is no longer an overwhelming endeavor.”

iWay Big Data Integrator can natively ingest, cleanse, and integrate your Hadoop-based data, making it work with your existing applications and processes so you can transform content into valuable insight. It reduces coding and maintenance, minimizes costs and complexity, and promotes best practices in data movement. It also works natively within any Hadoop environment, leverages native Hadoop capabilities and resource negotiation, and uses a wide variety of data sources and formats.

⁹ “Big Data,” Forrester Research, Inc. 2016.

Glossary

Apache Knox Gateway – Provides a single point of authentication and access for Apache™ Hadoop services.

Apache Sentry – A system for enforcing fine-grained role-based authorization to data and metadata stored on a Hadoop cluster.

Apache Software Foundation – A non-profit organization, which produces open-source software that powers much of the Internet behind the scenes.

Avro – A remote procedure call and data serialization framework developed within Apache's Hadoop project. It uses JSON for defining data types and protocols, and serializes data in a compact binary format. Its primary use is in Apache Hadoop, where it can provide both a serialization format for persistent data, and a wire format for communication between Hadoop nodes, and from client programs to the Hadoop services.

Cloudera – An American-based software company that provides Apache Hadoop-based software, support and services, and training to business customers. Cloudera's open-source Apache Hadoop distribution, CDH (Cloudera Distribution including Apache Hadoop), targets enterprise-class deployments of that technology. Cloudera says that more than 50 percent of its engineering output is donated upstream to the various Apache-licensed open source projects (Hive, Avro, HBase, etc.).

Cloudera Impala – Cloudera's open-source, massively parallel processing (MPP) SQL query engine for data stored in a computer cluster running Hadoop.

Data Lake – A method of storing data within a system that facilitates the collocation of data in variant schemata's and structural forms, usually object blobs or files. Hadoop, Azure Storage, and the Amazon S3 platform can be used to build data lake repositories.

Flume – A distributed, reliable, and available service for efficiently collecting, aggregating, and moving large amounts of log data. It has a simple and flexible architecture based on streaming data flows.

Hadoop – A set of open-source programs and procedures (for anyone to use or modify, with a few exceptions), which users can leverage as the backbone of their big data operations. Hadoop is made up of modules, each of which carries out a particular task essential for a computer system designed for big data analytics. Four of these modules are defined below. One additional module not formally part of Hadoop is Spark. Hadoop supports both traditional MapReduce and Spark for data processing.

- **Hadoop Distributed File System (HDFS)** – Allows data to be stored in an easily accessible format, across a large number of linked storage devices, and the MapReduce, which provides the basic tools for poking around in the data. (A file system is the method used by a computer to store data, so it can be found and used. Normally, this is determined by the computer's operating system, however a Hadoop system uses its own file system which sits "above" the file system of the host computer, meaning it can be accessed using any computer running any supported OS)

- **MapReduce** – Responsible for data/parallel processing. This module is named after the two basic operations this module carries out – reading data from the database, putting it into a format suitable for analysis (map), and performing mathematical operations, i.e., counting the number of males aged 30+ in a customer database (reduce). MapReduce is designed to run batch jobs that address every file in the system. Since that process takes time, MapReduce is well suited for large distributed data processing where fast performance is not an issue. Spark has become an alternative to MapReduce. Spark, unlike MapReduce, is all about performing sophisticated analytics at lightning fast speed in memory
- **Hadoop Common** – Provides the tools (in Java) needed for the user’s computer systems (Windows, Unix, etc.) to read data stored under the Hadoop file system
- **YARN** – Provides resource management and a central platform to deliver consistent operations, security, and data governance tools across Hadoop clusters

HBase – Open-source, non-relational, distributed database modeled after Google’s BigTable and is written in Java. It is developed as part of Apache Software Foundation’s Apache Hadoop project and runs on top of HDFS (Hadoop Distributed Filesystem).

HCatalog – A centralized metadata management and sharing service for Apache Hadoop. It allows for a unified view of all data in Hadoop clusters and allows diverse tools, including Pig and Hive, to process any data elements without needing to know physically where in the cluster the data is stored.

Hive – A data warehouse infrastructure built on top of Hadoop for providing data summarization, query, and analysis. Hive supports analysis of large datasets stored in Hadoop’s HDFS and compatible file systems. It provides an SQL-like language called HiveQL with schema on read and transparently converts queries to MapReduce and Spark jobs.

Hortonworks – A business computer software company based in Santa Clara, California. The company focuses on the development and support of Apache Hadoop, a framework that allows for the distributed processing of large data sets across clusters of computers.

JSON (JavaScript Object Notation) – A language independent data format that uses human-readable text to transmit data objects.

Kafka – A fast, scalable, durable, and fault-tolerant publish-subscribe messaging system. Kafka is often used in place of traditional message brokers because of its higher throughput, reliability, and replication. Kafka works in combination with Storm, HBase, and Spark for real-time analysis and rendering of streaming data.

Kerberos – A computer network authentication protocol that works on the basis of “tickets” to allow nodes communicating over a non-secure network to prove their identity to one another in a secure manner. Its designers aimed it primarily at a client-server model, and it provides mutual authentication – both the user and the server verify each other’s identity.

MapR – A San Jose, California-based enterprise software company that develops and sells Apache Hadoop-derived software. The company contributes to Hadoop projects like HBase, Pig, Hive, and ZooKeeper. MapR’s Hadoop distribution claims to provide full data protection, no single points of failure, improved performance, and dramatic ease of use advantages.

Maven – A build automation tool used primarily for Java projects. The word maven means “accumulator of knowledge” in Yiddish. Maven addresses two aspects of building software: first, it describes how software is built, and second, it describes its dependencies.

Oozie – A workflow processing system that lets users define a series of jobs written in multiple languages – such as MapReduce, Pig, and Hive – then intelligently link them to one another. Oozie allows users to specify, for example, that a particular query is only to be initiated after specified previous jobs on which it relies for data are completed.

Pig – A Hadoop-based language. It is relatively easy to learn and is adept at very deep, very long data pipelines (a limitation of SQL). Pig was originally developed at Yahoo Research for researchers to have an ad hoc way of creating and executing MapReduce jobs on very large data sets. Pig can execute its Hadoop jobs in MapReduce, Apache Tez, or Apache Spark.

Python – A widely used high-level, general-purpose, interpreted, dynamic programming language. Its design philosophy emphasizes code readability, and its syntax allows programmers to express concepts in fewer lines of code than possible in languages such as C++ or Java.

Scala – A general purpose programming language. It runs on the Java platform (Java virtual machine) and is compatible with existing Java programs. Apache Spark is built on Scala, thus users proficient in Scala helps them to dig into the source code when something does not work as expected. It is the language that Big Data Integrator generates.

SPT (Shortest Processing Time) – A data processing scheduling rule that ensures the job with the shortest processing time is scheduled first.

Spark – Another parallel processing engine in Hadoop ecosystem. It is a framework for performing general data analytics on distributed computing cluster. It provides in-memory computations for increased speed and data process over MapReduce. It runs on top of existing Hadoop cluster and access Hadoop data store (HDFS). It can also process-structured data in Hive and Streaming data from HDFS, Flume, Kafka, or Twitter. Apache Spark currently supports multiple programming languages, including Java, Scala, and Python. It is written in Scala.

Spark MLlib – Spark’s machine learning (ML) library. Its goal is to make practical machine learning scalable and easy. It consists of common learning algorithms and utilities, including classification, regression, clustering, collaborative filtering, dimensionality reduction, as well as lower-level optimization primitives and higher-level pipeline APIs.

Spark Streaming – Brings Spark's language-integrated API to stream processing, letting you write streaming jobs the same way you write batch jobs. It supports Java, Scala, and Python.

SparkSQL – A component on top of Spark Core that introduces a new data abstraction called DataFrames, which provides support for structured and semi-structured data. Spark SQL provides a domain-specific language to manipulate DataFrames in Scala, Java, or Python. It also provides SQL language support, with command-line interfaces and ODBC/JDBC server.

Sqoop – A connectivity tool for moving data from non-Hadoop data stores – such as relational databases and data warehouses – into Hadoop. It is SQL for Hadoop. It allows users to specify the target location inside of Hadoop and instruct Sqoop to move data from Oracle, Teradata, or other relational databases to the target.

Storm – A system for processing streaming data in real time. Storm adds reliable real-time data processing capabilities to Enterprise Hadoop. Storm on YARN is powerful for scenarios requiring real-time analytics, machine learning, and continuous monitoring of operations. Storm integrates with YARN via Apache Slider.

Worldwide Offices

Corporate Headquarters

Two Penn Plaza
New York, NY 10121-2898
(212) 736-4433
(800) 969-4636

United States

Atlanta, GA* (770) 395-9913
Boston, MA* (781) 224-7660
Channels (770) 677-9923
Charlotte, NC (980) 215-8416
Chicago, IL* (630) 971-6700
Cincinnati, OH* (513) 891-2338
Dallas, TX* (972) 398-4100
Denver, CO* (303) 770-4440
Detroit, MI* (248) 641-8820
Federal Systems, D.C.* (703) 276-9006
Florham Park, NJ (973) 593-0022
Houston, TX* (713) 952-4800
Los Angeles, CA* (310) 615-0735
Minneapolis, MN* (651) 602-9100
New York, NY* (212) 736-4433
Philadelphia, PA* (610) 940-0790
Pittsburgh, PA (412) 494-9699
San Jose, CA* (408) 453-7600
Seattle, WA (206) 624-9055
St. Louis, MO* (636) 519-1411, ext. 321
Washington, D.C.* (703) 276-9006

International

Australia*
Melbourne 61-3-9631-7900
Sydney 61-2-8223-0600
Austria Raffaisen Informatik Consulting GmbH
Wien 43-1-211-36-3344
Brazil
São Paulo 55-11-3372-0300
Canada
Calgary (403) 718-9828
Montreal* (514) 421-1555
Toronto* (416) 364-2760
Vancouver (604) 688-2499
China
Peacom, Inc.
Fuzhou 86-15-8800-93995
SolventoSOFT Technology (HK) Limited
Hong Kong 852-9802-4757
Czech Republic InfoBuild Software CE s.r.o.
Praha 420-234-234-773
Estonia InfoBuild Estonia OÜ
Tallinn 372-618-1585
Finland InfoBuild Oy
Espoo 358-207-580-840
France*
Suresnes +33 (0)1-49-00-66-00
Germany
Eschborn* 49-6196-775-76-0
Greece Applied Science Ltd.
Athens 30-210-699-8225
Guatemala IDS de Centroamerica
Guatemala City (502) 2412-4212
India* InfoBuild India
Chennai 91-44-42177082
Israel SRL Software Products Ltd.
Petah-Tikva 972-3-9787273
Italy
Agrate Brianza 39-039-59-66-200
Japan KK Ashisuto
Tokyo 81-3-5276-5863
Latvia InfoBuild Lithuania, UAB
Vilnius 370-5-268-3327
Lithuania InfoBuild Lithuania, UAB
Vilnius 370-5-268-3327
Mexico
Mexico City 52-55-5062-0660

Netherlands*

Amstelveen 31 (0)20-4563333

- Belgium
- Luxembourg

Nigeria InfoBuild Nigeria
Garki-Abuja 234-9-290-2621

Norway InfoBuild Norge AS c/o Okonor
Tynset 358-0-207-580-840

Portugal

Lisboa 351-217-217-400

Singapore Automatic Identification Technology Ltd.
Singapore 65-69080191/92

South Africa InfoBuild (Pty) Ltd.
Johannesburg 27-11-064-5668

South Korea

Dfocus Co., Ltd.
Seoul 02-3452-3900
UVANSYS, Inc.
Seoul 82-2-832-0705

Southeast Asia Information Builders SEAsia Pte. Ltd.
Singapore 60-172980912

- Bangladesh ■ Brunei ■ Burma ■ Cambodia
- Indonesia ■ Malaysia ■ Papua New Guinea
- Thailand ■ The Philippines ■ Vietnam

Spain

Barcelona 34-93-452-63-85
Bilbao 34-94-400-88-05
Madrid* 34-91-710-22-75

Sweden

Stockholm 46-8-76-46-000

Switzerland

Brugg 41-44-839-49-49

Taiwan

Azion Corporation
Taipei 886-2-2356-3996
Galaxy Software Services, Inc.
Taipei 886-2-2586-7890, ext. 114

Tunisia North Africa Business Consulting
Bizerte 215-22-86-85-79

Turkey Paladin Bilisim Danismanlik A.S.
Istanbul 90-532-111-32-82

United Kingdom*

Uxbridge Middlesex 44-20-7107-4000

Venezuela InfoServices Consulting
Caracas 58212-763-1653

West Africa InfoBuild FSA
Abidjan 225-01-17-61-15

* Training facilities are located at these offices.

Find Out More

We can help you succeed. Talk to your local Information Builders representative to learn how. Visit us at informationbuilders.com, e-mail askinfo@informationbuilders.com, or call (800) 969-4636 in the U.S. and Canada. To improve your skills, visit education.ibi.com.

