

Cisco Information Server 8



“At Pfizer, when senior execs come to me with key questions whose answers will determine the courses of action we’ll take in running our business, my team uses the rapid methods built within Cisco Information Server. This reduces each project from 4-6 weeks to 2-3 days.”

—Dr. Michael Linhares, Director of Business Operations, Pfizer

With data the new competitive battleground, businesses that take advantage of their data will be the leaders; those that do not will fall behind.

But gaining this advantage is a more difficult technical challenge than ever as your business requirements are ever-changing, your analytic workloads are exploding and your data is more-widely distributed across on-premises, big data, the internet of things and the cloud. Traditional physical data integration via data warehousing and ETL is often too time consuming, too rigid and too costly to support your dynamic business needs. There has to be a better way. There is and it is called data virtualization.

Cisco® Information Server is data virtualization software that lets you integrate data at big data scale, with breakthrough speed and cost effectiveness. With Cisco Information Server, you can build and manage virtualized views and data services that access, transform and deliver the data your business requires to accelerate revenue, reduce costs, lessen risk, improve compliance and more. Cisco Information Server is:

- **Fast and Economical** – Integrate data reliably at a fraction of physical warehousing and ETL’s time, cost and rigidity. Evolve rapidly when requirements change.
- **Immediate** – Deliver up-to-the-minute data as needed, using advanced performance optimization algorithms, massively parallel processing and fine-grained security.
- **Business-friendly** – Transform native IT structures and syntax into easy-to-understand, IT-curated data sets sharable via a self-service business directory.
- **Wide-ranging** – Access data from distributed data sources including traditional enterprise, big data, cloud and IoT. Use it across myriad analytics, self-service, business intelligence and transactional applications.
- **Enterprise Grade** – Support multiple lines of business, hundreds of projects and thousands of users.

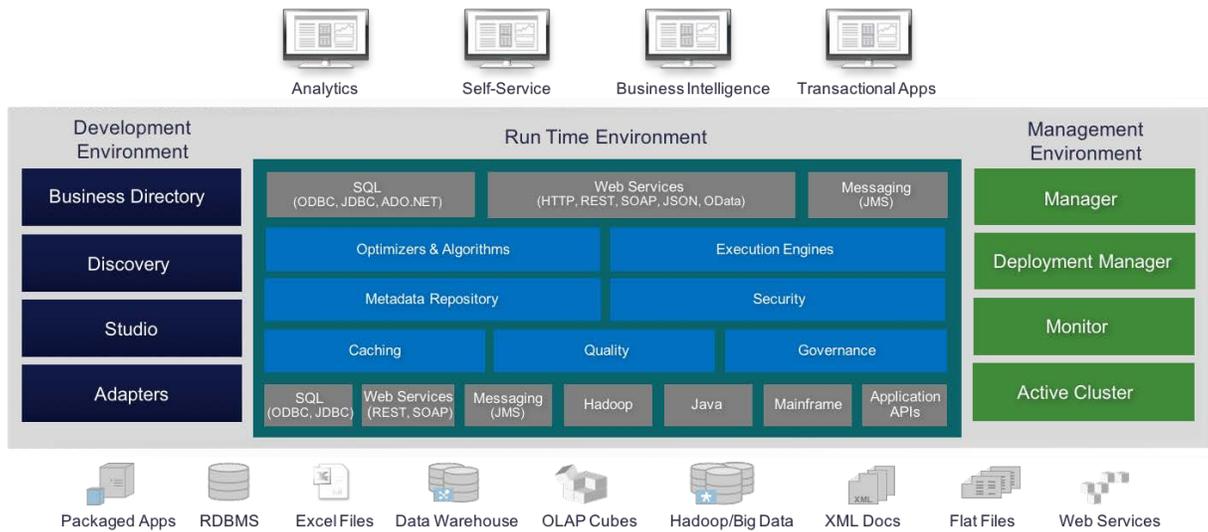


Figure 1. Cisco Information Server Modules

Cisco Information Server Modules

The Cisco Information Server is Java-based, enterprise-grade middleware whose modular structure supports all phases of data virtualization development, run-time and management.

- **Business Directory** is a self-service business data directory you can use to easily search, categorize and consume IT-curated data sets developed using the Cisco Information Server.
- **Discovery** enables you to go beyond simple data profiling to examine data, locate important entities and reveal hidden relationships across distinct data sources. You can quickly build and display comprehensive E-R diagrams and data models so you can meet new business requirements faster and more easily.
- **Studio** is the agile modeling, development and resource management tool your data-oriented developers use to model data, design views and data services, build transformations, optimize queries, manage resources, and more. Easy to learn and use, Studio's graphical modeling environment along with powerful code editors provide a flexible workspace where queries are created and tested, as well as a data services repository where completed objects are published. Studio also offers a rich set of transformations in addition to an easy-to-use transformation editor. Five development languages including SQL, SQL Script, Java, XQuery and XSLT complement Studio's graphical modeling capabilities.
- **Adapters** provide a wide range of data source connectivity for databases, files, big data, cloud sources, packaged applications and more. Beyond schema-to-schema only connectivity, CIS Adapters integrate CIS optimizers with data source optimizers to ensure more accurate queries and higher performance. The **Data Source Tool Kit** allows you to easily build additional custom adapters.
- **Optimizers and Algorithms** intelligently evaluate workloads at runtime and dynamically select the ideal combination of optimizers (cost-based, rule-based and data sources), join algorithms, query techniques and execution engines (CIS's new MPP Execution Engine or CIS's Classic Engine) to achieve the best possible performance.

- Multiple **Execution Engines** allow you to run optimized data virtualization workloads at any scale. CIS's Classic Execution Engine processes workloads on a single node within a clustered environment. CIS's Massively Parallel Processing Execution Engine (MPP Execution Engine) dynamically distributes big data scale integration workloads across processing nodes.
- The integrated **Metadata Repository** lets you store and manage all relevant metadata including your views and data services throughout their development and operations lifecycle.
- Myriad fine-grained **Security** capabilities, including authentication, authorization encryption and masking, leverage your existing security investments (LDAP, Active Directory and Kerberos), standards (SAM, NTLM, TLS and WSS) and policies to securely access and deliver sanctioned data only, no more, no less.
- Flexible **Caching** options enable higher performance and greater uptime by efficiently moving data to predesignated caches.
- **Quality** provides a range of profiling, standardization, conformance, enrichment, augmentation and validation capabilities that you can apply to ensure delivery of the most correct and complete data possible.
- Built-in **Governance** features including lineage, where-used, logging and provide complete visibility, traceability and control over your data virtualization activities.
- **Manager** is the administrative console your administrators use to set up user IDs, passwords, security profiles, etc., as well as check server activity, and more.
- **Deployment Manager** lets you quickly and easily migrate entire projects in a single step, including their resources, cache settings, security profiles and more across multiple CIS instances to simplify and automate your development lifecycle.
- **Monitor** provides a comprehensive, real-time view of your CIS clusters. Monitor displays all the pertinent system health indicators required to help your IT operations staff guide corrective actions.
- **Active Cluster** works in conjunction with your load balancers to provide high availability and greater scale to meet your challenging service-level-agreements. Active Cluster simplifies complex operations management by automatically sharing resources, adjusting capacity on demand, and more.

Deployment Options

Cisco provides multiple options for deploying Cisco Information Server. You can install and run CIS on-premises, in your private cloud environment or at a public cloud provider such as Amazon AWS, Google Cloud Platform and Microsoft Azure. For Amazon, Cisco also provides an Amazon Machine Image (AMI) on the AWS Marketplace to simplify and accelerate deployment.

Complete Data Virtualization Solution

Beyond the Cisco Information Server, Cisco also provides a broad array of complementary **Advanced Services, Training, Technical Support, Customer Advisory Program, Knowledgebase** and **Partner** offerings so you get the complete solution needed to ensure your data virtualization success.

Cisco Information Server 8 at a Glance

Ten times faster than CIS 7, Cisco Information Server 8 (CIS 8) sets the standard for big data scale data virtualization with unique and unprecedented advancements in data virtualization scalability, performance and workload management.

Massively Parallel Processing Execution Engine Powers Big Data Scale Virtualization

Until CIS 8, big data scale integration workloads required dedicated storage and processing infrastructure and all the cost, time and rigidity associated with this approach. With CIS 8's Massively Parallel Processing Execution Engine, you no longer need to compromise when providing the analytic data your business requires. With CIS 8 you can now deliver:

- **Big Data Scale Integration Workloads without Big Data Repositories** – CIS 8's new massively parallel processing execution engine dynamically distributes your big data scale integration workloads across processing nodes. You get the performance and scale you need, without the extra cost and management headaches associated with replicated data and dedicated storage infrastructure.
- **Integrated Data Delivery without Sacrificing Agility and Cost-Effectiveness** – With CIS 8, you no longer need to trade-off data virtualization's agility and cost savings for runtime performance. CIS 8 is the only enterprise-grade data virtualization offering that lets you quickly and cost-effectively meet business demand for integrated data at big data scale.
- **10 Times Faster Data Virtualization without Extra Engineering** – With CIS 8 you don't need to reengineer your views, data services and IT infrastructure to achieve the performance you need. CIS 8's intelligent optimizers, algorithms and engines automatically optimize your workloads and resources to maximize performance for you.

Advanced Workload Management for Big Data Scale SLAs

CIS 8 also helps you achieve your SLAs at big data scale. Advancements in CIS 8 workload management provide more granular resource controls that help ensure your most important workloads get prioritized. With CIS 8 you can now control memory usage, request length, row counts, and more, as well as avoid potentially problematic requests, for example queries without filters, large cross joins, etc. You can implement these controls at the object, user group, server, and / or cluster group level and monitor them using existing CIS reporting.

New Adapters, for Big Data and More, Turn Your Data Silos into an Advantage

CIS 8 lets you turn your big data silos into an advantage. Where big data repositories are required, CIS 8 can easily access them using new data source adapters for Apache Hive 2.0, Apache Spark SQL 2.1, Cloudera Hive 2.0 and Cloudera Impala 2.0, to now deliver nearly twenty Hadoop, NoSQL and cloud big data sources. CIS 8 also provides an Elastic Search adapter for the first time. And as with every CIS release, multiple others adapters have been updated to ensure the accuracy and performance you require when accessing your diverse, distributed sources.

Unleash Big Data Scale Virtualization with Developer Productivity Advancements

To help you take advantage of CIS's unprecedented performance, CIS 8 raises the bar on developer ease-of-use and time-to-solution agility. Common development tasks such as auto-completion and expanding / collapsing code blocks and more are easier with Studio's expanded code editors. Plus, additional JSON formatting functions further simplify and accelerate REST services development.

Key Features

Development Environment

Tables 1 through 4 summarize key Cisco Information Server development capabilities.

Table 1. Business Directory: Access and Use IT-Curated Data Sets via an Intuitive, Self-Service Browser

Feature	Description
Search options	Find data sets using simple Google-like searches as well as advanced parameterized searches.
Browse	Find data sets using attributes such as data type, format, category and more.
Data preview	Preview specific data included within data sets.
Security	See only the data sets you're supposed to see based on your Cisco Information Server enforced security profile.
Consumption	Use your favorite analytic/BI tools to consume data sets using Business Directory access information.
Categorization	Organize large and / or diverse data sets into system or user-defined categories.
Collaboration	Business users and IT can interact via comments to improve data quality and utilization.
Business metadata extensions	Add custom definitions, properties, links and status codes to enrich IT metadata with Business metadata your business users understand.
Business metadata display options	Present business metadata in a shared area, an existing properties tab or a new custom tab.
Personalization	Receive email updates of changes and comments for data you care about.
Provisioning	Register additional Cisco Information Server instances and expand available data sets.
REST API	Access Business Directory programmatically to expose data sets to additional consuming applications.

Table 2. Modeling and Transformation: Design Reusable Views and Data Services within an Agile, High-Productivity Development Environment

Feature	Description
Graphical development environment	Graphically model data, design views and data services, build transformations, optimize queries, manage resources, and more.
Agile development methodology	Develop and refine views and data services iteratively.
Intelligent code editors	Myriad code editors for common tasks such as auto-completion, expanding / collapsing code blocks and more improve developer productivity.
Introspection	Automatically probe physical data sources and select desired resources. Optionally inspect data sources interactively.
Data discovery	Automatically determine data relationship across disparate entities using formal keys and fuzzy matching.
E-R diagrams	Visualize discovered relationships and data models via E-R Diagrams.
Bottom-up modeling	Design views and data services by combining data from disparate systems without worrying about underlying source access and format complexity.
Flexible modeling and transformation methods	Create views and data services graphically or via scripting languages as appropriate including SQL, SQL Script, XQuery, XSLT and Java functions.

SQL	Use familiar SQL standards.
SQL Script	Implement stored procedures using a familiar scripting language.
XQuery	Create complex XML structures using a graphical XQuery editor.
XPath transformation	Establish arbitrary complex mapping of XML schema elements to XML output.
JSON querying and transformation	Query and transform JSON data from Web services into a relational format and vice versa.
Analytic functions	Use a full set of analytic functions such as CORR, COUNT, NTILE, STDDEV, VARIANCE, etc. pushed to the data source or executed by CIS as appropriate
Contract first design	Build data services using preexisting WSDLs and schemas.
Contract last design	Define the data integration solutions first and then wrap them with a WSDL.
Query plan	View query execution plan steps and details in Studio or through client drivers.
Views dependency graph	Graphically display dependencies between data sources and views and data services.
Physical table creation	Create and drop physical tables within a designated data source.
Localized UI	Chinese and Japanese language support.

Table 3. Metadata Management: Store and Manage Relevant Metadata

Feature	Description
Complete repository	Manage resources such as data sources, views, data services and procedures throughout their lifecycles.
Public metadata API	Deploy Web Services–based metadata API for easy metadata access and sharing.
Schema change notification	Receive notice when data source schemas change.
Source metadata	Access and use metadata from physical data sources.
Open API	Open Cisco Information Server system libraries access to enable custom scripting and orchestration.

Table 4. Integrated Version Control: Control the View / Data Service Development Lifecycle

Feature	Description
GUI-based migration	Add project folders directly to version control systems. Check-in and check-out folders or individual resources to track changes. Rollback revisions of folders or resources to prior versions.
Resource locking	Protect against inadvertent modifications and overwrites.
Change history	Track changes made by users with annotations.
Source control and transport	Manage artifacts and transports from development through production.
Version control systems	Integrate directly with GIT, SVN and TFS.

Run-Time Environment

Tables 5 through 12 summarize key Cisco Information Server run-time operations capabilities.

Table 5. Intelligent Optimization Algorithms and Techniques: Automatically Optimize Query Plans

Feature	Description
Intelligent query optimization	Intelligently evaluate workloads at runtime and dynamically select the best combination of optimizers (cost-based, rule-based and data sources), join algorithms, query techniques and execution engines (classic and massively parallel).
Cost-based optimizer	Use query characteristics, table statistics, indexes and data distribution histograms to create an optimal query plan.
Rule-based optimizer	Specify rules that can influence a query plan.
Data source optimizers	Leverage data source optimizers to ensure query accuracy and maximize query performance.
Complete set of join algorithms	Join and aggregate vertically and horizontally partitioned data using full range of join algorithms including hash join, sort-merge join, distributed semi-join, data-ship join, union-join flip, nested-loop join, etc.
Automatic query rewrite	Automatically rewrite the query to use the most efficient join strategy.
SQL pushdown	CIS offloads as much processing as possible by pushing down select query operations such as string searches, comparisons, local joins, sorting, aggregating and grouping into the underlying data sources.
Single-source join grouping	Run data-reducing joins at the data source rather than bringing the data across the network.
Predicate push-down	Push WHERE clause predicates down into the underlying data source to reduce data at the source.
Full and partial aggregate push down	Push aggregate functions down to source when applicable
Serialization or parallelization of join operators	Determine the proper join order and join algorithms based on estimated cardinality and join results derived from data distribution histograms.
Projection pruning	Eliminate all unnecessary columns from fetch nodes in a query tree.
Constraint propagation	Distribute filters to multiple branches of the query plan, allowing data reduction by a single filter to potentially occur in multiple data sources.
Scan multiplexing	Reuse data sets that appear in multiple places in a single query plan.
Empty scan detection	Detect logical conditions that would produce empty data sets, and then eliminate those parts of the query plan prior to processing.
Redundant operator cropping	Eliminate redundant or extraneous operators within a complex multiple-operator query.
Blocking operator prefetching	Proactively run parts of the query plan that must finish before other parts of the query plan can continue, thereby increasing the overall responsiveness of the query.
Native XML support	Support XML internally for faster parsing and joins.
Caching	Store pre-queried result sets either in a local file-based cache or in a database.
API	Expose query execution plan via JDBC/ODBC.

Table 6. Multiple Execution Engines: Run Optimized Data Virtualization Workloads at any Scale

Feature	Description
Dual execution engines	Automatically run the classic execution or the massively parallel execution engine based on the optimizer's intelligent evaluation.
Classic execution engine	Process workloads on a single node within a clustered environment.
Massively parallel processing execution engine (MPE)	Dynamically distribute big data scale integration workloads across nodes to maximize performance at scale.

Multi-threaded, asynchronous access	Execute calls asynchronously on separate threads, reducing idle time and data source response latency.
Connection pool sharing	Share access to data sources to avoid bottlenecks.
Hybrid memory and disk usage	Automatically balance memory and disk use for optimal performance.
Fine-grained workload management	Monitor and conditionally implement rules based on server workload thresholds. When rules and thresholds are violated, execute appropriate actions (Run with Administrator alert, Run with User warning, Stop execution with Administrator alert and User exceptions).
Results streaming	Stream data to consuming applications as results are processed at the underlying sources.
Scheduling	Run queries based on set times.
Alert triggers	Implement resource, event and user-defined triggers.
Alerts API	Use a published API to handle custom Java alerts.

Table 7. Flexible Caching: Move Data to Predesignated Caches to Boost Availability and Performance

Feature	Description
Event-based refresh	Refresh caches based on defined business rules.
Scheduled refresh	Refresh caches at set times.
Manual refresh	Refresh caches on demand as needed.
Incremental refresh	Refresh part of a cache.
Full refresh	Refresh entire cache.
Native bulk extract	Use native bulk EXTRACT functions in source to extract data more efficiently than using SELECTs.
Native bulk load	Use native LOAD functions in destination to load and refresh caches most efficiently than using INSERTs.
Parallel load	Use multiple threads to load caches in parallel.
Multi-table caching	Avoid contention on cache refreshes, accelerate refresh speeds and maintain non-stop availability using multiple tables per cache view.
Centralized cache tracking	Track caching tables distributed across multiple data sources centrally.

Table 8. Data Access: Connect and Use Data from Distributed Sources

Feature	Description
Relational data	Access popular databases via Java Database Connectivity (JDBC).
Hadoop data	Access Hadoop sources through Drill, Hive, HBase, Impala, SparkSQL and more.
Multidimensional data	Access multidimensional data in SAP BW.
Cloud data	Access cloud database sources such as Amazon RedShift and Google BigQuery.
NoSQL data	Access NoSQL sources such as Cassandra, Couchbase and MongoDB.
Flat files	Access flat files (cache, delimited and XML).
Binary data	Access BLOB and CLOB data.
SOAP Web Services	Access SOAP over HTTP (XML).
REST Web Services	Access REST over HTTP (JSON or XML).
OData	Access data via OData protocol.
Java Message Service (JMS)	Access data via Java Message Service (JMS) services pipeline.
Elastic Search	Access data via elastic search.
Packaged applications	Access SAP, Oracle E-Business Suite, Salesforce.com and other applications using standard objects such as invoices, shipments, orders, customers, opportunities and more.
CRM and ERP SAAS applications	Access CRM and ERP SAAS applications Google Analytics, Microsoft Dynamics and NetSuite.

Collaboration applications	Access collaboration applications such as Email, Google Spreadsheets and Microsoft SharePoint.
Social media applications	Access social media sources such as Facebook, LinkedIn, RSS and Twitter.
Streaming / IoT data	Access streaming and IoT data via Cisco Edge Fog Fabric.
Data Source Tool Kit	Access a set of libraries of services that can be imported into your preferred Integrated development kit to facilitate and accelerate custom data adapter creation. Services include database mapping, data type mapping, syntax mapping, and function mapping. Together these services minimize custom code development.

Table 9. Data Delivery: Deliver Requested Data to Myriad Consuming Applications

Feature	Description
Database objects	Publish views for consumption through ODBC, JDBC and ADO.NET.
Web services	Publish data services in the form of WSDL for consumption using SOAP or SOAP over JMS. A message pipeline allows interjection of custom logic during the web service request and response.
Representational State Transfer (REST)	Publish data services in the REST format. REST CREATE, READ, UPDATE and DELETE functions are supported.
Open Data (OData) protocol	Publish data services in the OData format.
JSON	Publish JSON, including more formatting capabilities for XML-to-JSON translation, for both procedures and table outputs.
BI Tool Integration	BI tools can create temp tables to store filters for visualizations or reports.
DDL support for client applications	Support for client applications to create tables using DDL statements via JDBC/ODBC/ADO. CREATE TABLE, CREATE TABLE AS SELECT and CREATE TEMP TABLE syntax. Multiple physical data sources mappings are supported for each published database to avoid conflicts between users when creating tables.

Table 10. Integrated Security: Protect Sensitive Data with Fine-grained, Multi-level Security

Feature	Description
Policy-based security	Apply authentication, authorization and encryption rules via policies.
Single sign-on	Sign on once to access CIS, integrated data sources and consuming applications.
Authentication types	Basic, Pass-through, Kerberos, SAML and NTLM
Row-level authorization	Control access to specific rows via granular permissions.
Column-level authorization	Control access to specific columns via granular permissions.
Column masking	Implement column masking rules to hide, replace or obfuscate portions of a column's value depending on a user's level of access.
SSL over HTTP with support for mutual authentication	Mutually authenticate published services, web services data sources and Oracle databases. Certificate-based authentication and Web Services Security (WSS) authentication are supported.
Pass-through	Use an existing user ID and password and pass through to Cisco Information Server for authentication.
Lightweight Directory Access Protocol (LDAP)	Use security profiles from LDAP or Active Directory to authenticate user access to protected data sources.
Pluggable authentication module	Use third-party systems for authentication.
Encryption	Encrypt passwords and data in motion via TLS.
Access management	Use Cisco Information Server as the system of record for security roles and profiles.

Table 11. Data Quality: Help Ensure Correct and Complete Data

Feature	Description
Profiling	Introspects data sources to profile structure, statistics and content. Generates full data profile (counts, duplicates, min, max, histograms, etc.) and detects errors, outliers and duplicates.
Reconciliation	Automatically reconcile non-matching data elements such as case-sensitivity and trailing spaces.

Standardization and conformance	Create views and data services that conform to agreed standards.
Enrichment and augmentation	Extend views and data services with additional data.
Validation	Validate data sets with users prior to publishing views and data services.
Object reuse	Share views and data services to ensure consistent data definitions.

Table 12. Governance: Provide Comprehensive Visibility, Traceability and Control

Feature	Description
Data lineage	Trace lineage from multiple data sources to a single data consumer.
Column lineage API	Access column lineage via an API.
Where-used	Trace where-used from a single data source to multiple data consumers.
Logging	Track system and/or user activity.
Standards enforcement	Implement internal and industry data standards.
Open API	Open access to Cisco Information Server system libraries to enable custom scripting and orchestration.

Management Environment

Tables 13 through 15 summarize key Cisco Information Server management capabilities.

Table 13. User and System Management: Administer and Manage for Efficient Operations

Feature	Description
Management console options	Access the management console through Studio or a Web browser.
User setup	Set up user and group profiles.
Security setup	Enable multiple forms of security.
Scheduling	Schedule Cisco Information Server run-time activities flexibly.
Task management	Manage tasks related to development, management, configuration and versioning.
Workload management setup	Define workload management rules and server management thresholds (memory usage, request length, row counts, and more) at the object, user group, server, and / or cluster node level.
Fine-grained workload management	Monitor and conditionally implement rules based on server workload thresholds. When rules and thresholds are violated, execute appropriate actions (Run with Administrator alert, Run with User warning, Stop execution with Administrator alert and User exceptions)
Real-time system indicators	Track critical system metrics and tune for optimal performance including memory usage, query plans for (currently running and past requests), data sources and caches.
Usage metrics	Deliver usage activity detail to your reporting tool of choice via an open API.
Simple Network Management Protocol	Allow monitoring by third-party systems via SNMP API.

Table 14. Cluster Management: Cluster Servers to Meet Reliability, Availability and Scalability Service-Level-Agreements

Feature	Description
Flexible cluster deployment	Expand capacity and improve availability simply by adding new nodes to an existing cluster or adding new clusters.
Active / Active clustering	Distribute Cisco Information Server workloads across cluster nodes in conjunction with load balancers.
Shared cluster cache	Improve overall cluster performance by coalescing redundant data source hits and reducing data latency.
Automated metadata synchronization	Replicate metadata across clusters to simplify management.
Restore cluster nodes	Restore cluster nodes using Web Manager, cluster_util script or API.

Table 15. Deployment Management: Automate Migration / Promotion of Artifacts, Configurations and Settings

Feature	Description
Resource migration	Migrate / promote (create/update/delete) artifacts from one Cisco Information Server instance to another.
Cache setting migration	Migrate / promote cache table names, caching methods, refresh method, cache policies and cache schedules.
Trigger migration	Migrate / promote triggers.
User/group migration	Migrate / promote user / group IDs, security profiles and other user / group information.
Remote Server API Calls	Execute remote API calls on CIS instances to execute custom operations such as update server configurations (e.g. enabling/disabling triggers)

Technical Specifications

Tables 16 through 18 summarize Cisco Information Server's technical specifications and standards support.

Table 16. Supported Platforms: Leverage Standard Operating Systems and High Performance Servers

<p><u>CIS Studio Clients</u></p> <ul style="list-style-type: none"> • Microsoft Windows <p><u>Client Drivers</u></p> <ul style="list-style-type: none"> • ADO.NET • ODBC for Windows Linux, AIX, HP-UX and Solaris • JDBC 	<p><u>Server</u></p> <ul style="list-style-type: none"> • CentOS • Cisco Unified Computing System™ • JRE 1.8 • Oracle Linux and Red Hat compatibility mode • Red Hat Enterprise Linux AS • SUSE Enterprise Linux • Microsoft Windows • Solaris 	<p><u>JVM</u></p> <ul style="list-style-type: none"> • 64-bit
---	---	---

Table 17. Data Source Adapters: Simplify Data Access while Ensuring Query Accuracy and Performance

<p><u>Relational Adapters</u></p> <ul style="list-style-type: none"> • Cisco Information Server • Data Direct Mainframe • Greenplum • HSQLDB • IBM DB2 • IBM DB2 z/OS • Informix • Microsoft Access • Microsoft Excel • Microsoft SQL Server • MySQL • Netezza NPS • Oracle • Parstream • PostgreSQL • SAP HANA • Sybase • Sybase IQ • Teradata • Vertica <p><u>Web Services and File Adapters</u></p> <ul style="list-style-type: none"> • Custom Java Procedure • Elastic Search • Flat Files • LDAP / Active Directory • OData • REST • SOAP • XML <p><u>Cloud Database Adapters</u></p> <ul style="list-style-type: none"> • Amazon RedShift • Google BigQuery 	<p><u>Hadoop Adapters</u></p> <ul style="list-style-type: none"> • Apache Drill • Apache HBase • Apache Hive • Apache Impala • Apache Spark SQL • Cloudera Hive • Cloudera Impala • Hortonworks HDP <p><u>In-Memory Database Adapters</u></p> <ul style="list-style-type: none"> • HSQLDB • IBM DB2 In-Memory Option • Microsoft SQL Server In-Memory Option • Oracle In-Memory Option • SAP HANA <p><u>NoSQL Adapters</u></p> <ul style="list-style-type: none"> • Amazon DynamoDB • Cassandra • Couchbase • MongoDB <p><u>Collaboration Application Adapters</u></p> <ul style="list-style-type: none"> • Email • Google Apps • Google Sheets • Microsoft Active Directory • Microsoft SharePoint (On-premise and online) • Microsoft SharePoint Excel Services <p><u>Data Source Tool Kit</u></p> <ul style="list-style-type: none"> • Software Development Kit for Custom Data Source Adapter Development 	<p><u>CRM and ERP Applications Adapters</u></p> <ul style="list-style-type: none"> • Microsoft Dynamics CRM (On-premise & Online) • Microsoft Dynamics GP • Microsoft Dynamics NAV • NetSuite CRM • NetSuite ERP • Oracle EBS • Salesforce.com • Siebel <p><u>Marketing Automation Applications Adapters</u></p> <ul style="list-style-type: none"> • Google Adwords • Google Analytics • HubSpot • Marketo • Oracle Eloqua <p><u>SAP Adapters</u></p> <ul style="list-style-type: none"> • SAP Netweaver BW • mySAP Business Suite • SAP Business Explorer (BEx) • SAP HANA <p><u>Social Media Adapters</u></p> <ul style="list-style-type: none"> • Facebook • LinkedIn • RSS • Twitter <p><u>Streaming / IoT Data Adapters</u></p> <ul style="list-style-type: none"> • Cisco Edge Fog Fabric
---	--	---

Table 18. Additional Technical Specifications: Take Advantage of Industry Standards

<p><u>Delivery Interfaces</u></p> <ul style="list-style-type: none"> • ADO.NET • ODBC • Hadoop • JDBC • REST • SOAP over HTTP • SOAP over JMS <p><u>Enterprise Service Buses</u></p> <ul style="list-style-type: none"> • OpenMQ • Sonic MQ • TIBCO EMS <p><u>Web Services Protocols</u></p> <ul style="list-style-type: none"> • .NET • OData • REST and JSON • SOAP • WSDL • WSI • XPath • XQuery • XSLT • XML (flat files or over HTTP) 	<p><u>Security Protocols</u></p> <ul style="list-style-type: none"> • Base64 • Kerberos • NTLM • SAML • SSL • WS-Security <p><u>Directory Services</u></p> <ul style="list-style-type: none"> • Active Directory • Oracle Directory Server Enterprise Edition • Cisco Information Server • Novell eDirectory <p><u>Software Development Standards</u></p> <ul style="list-style-type: none"> • SQL 92 and 99 • Unicode • JDK • J2EE • JNDI 	<p><u>Cache Repositories</u></p> <ul style="list-style-type: none"> • File • Greenplum • HSQLDB • IBM DB2 • Microsoft SQL Server • MySQL • Netezza • Oracle • PostgreSQL • SAP HANA • Sybase ASE • Sybase IQ • Teradata • Vertica <p><u>Data Ship Sources and Targets</u></p> <ul style="list-style-type: none"> • IBM DB2 • Microsoft SQL Server • Netezza • Oracle • PostgreSQL • Sybase IQ • Teradata • Vertica
---	--	--



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)