

Availability and Performance for Informatica PowerCenter and B2B Data Exchange

InfoScale Enterprise

Contents

Executive Summary	3
Solution Value	3
Solution Components	3
Availability	5
Maximum Availability for PowerCenter and B2BDX	5
Intelligent Monitoring	8
Performance and Optimization	9
Cluster File System (CFS)	9
Flexible Storage Sharing (FSS)	11
I/O Fencing	11
Efficiency and Agility	12
Veritas Volume Replication and Global Cluster Option	12
Veritas InfoScale Operations Manager	12
Public Cloud Support	13
Summary	14
Table of Figures	15

Executive Summary

With the increasing importance of Enterprise Data Integration (EDI), Informatica has become a leading solution for managing enterprise data transformation and integration initiatives. Informatica PowerCenter and B2B Data Exchange (B2BDX) are leading solutions for enterprise data warehousing, integration, transformation and data exchange. With a primary focus on providing a feature-rich platform, there are some challenges to consider when building an Informatica EDI solution:

- Meeting uptime requirements for the Informatica application and the underlying hardware
- Delivering the high performance and resiliency required for the underlying data storage system

Veritas InfoScale Enterprise is the solution recommended by Informatica to manage high availability and high-performance storage for PowerCenter and B2BDX. With over 30 years of innovation and market leadership in both Software-Defined Storage, application-aware resiliency and disaster recovery, InfoScale Enterprise provides intelligent volume management, an advanced high-performance shared file system and application-aware system clustering for PowerCenter and B2BDX. InfoScale Enterprise enables Informatica customers to rapidly deploy highly available, performant, multi-tier Informatica EDI services on nearly any type of infrastructure, operating system, platform, or underlying storage and compute infrastructure.

Solution Value

InfoScale Enterprise has customized integration with both PowerCenter and B2BDX and provides the availability, resiliency and performance required to satisfy enterprise usability expectations and SLA's for Informatica EDI environments. Using InfoScale Enterprise with Informatica results in a highly available and resilient platform for managing enterprise data transformation initiatives.

InfoScale Enterprise provides the following key benefits for Informatica:

- ✓ Intelligent monitoring of application health and proactive instant failure detection
- ✓ Agent-based monitoring and automation that ensures maximum PowerCenter and B2BDX availability
- ✓ Scalable and fault-tolerant clustered file system for Informatica HA deployments
- ✓ Reduced storage costs with a flexible scale-out architecture
- ✓ Data integrity in the event of a node failure in PowerCenter HA and B2BDX HA deployments
- ✓ High availability and shared storage for PowerCenter and B2BDX in public cloud environments
- ✓ Intelligent replication for agility and mobility across geographical locations
- ✓ Automated high availability management for multi-tier integrated Informatica services

Overall, InfoScale Enterprise helps maintain maximum availability for Informatica EDI projects and provides a high performance and resilient shared storage environment recommended by Informatica for HA deployments.

Solution Components

The following are key components in the InfoScale Enterprise solution for Informatica:

- **Cluster File System (CFS)** – An NFS server failure can cause downtime for PowerCenter and B2BDX servers. CFS solves this problem and improves operational efficiency by eliminating the need for separate NFS storage systems (primary and standby) to maintain high availability. CFS is recommended by Informatica for PowerCenter HA and B2BDX HA. CFS provides highly available, parallel access to files across all nodes in an Informatica HA configuration. CFS provides better lock management over NFS and makes failover more reliable.
- **Flexible Storage Sharing (FSS)** – Storage costs for Informatica are one of the leading factors affecting the Total Cost of Ownership (TCO). Often expensive SAN storage is typically used to meet the shared storage requirement for HA deployments. InfoScale FSS is a feature of CFS that provides a more cost effective and higher performing shared storage solution than SANs. FSS allows network shared storage to co-exist with physically shared storage, and logical volumes to be created using both types of storage, enabling a common storage namespace. Logical volumes using network shared storage provide data redundancy, high availability, and disaster recovery capabilities, without requiring physically shared storage. FSS is transparent to file systems and applications. FSS can be implemented using Direct Attached Storage (DAS) on PowerCenter and B2BDX nodes.
- **Veritas Volume Replicator (VVR)** – Enables platform independent disaster recovery by intelligently managing the replication of Informatica data between sites, regardless of the underlying infrastructure. VVR can replicate Informatica data between on-premises data centers, from an on-premises data center to a public cloud environment or between different public cloud environments. When integrated with the InfoScale Global Cluster Option, VVR provides optimized Informatica data replication between geographically dispersed sites.
- **Veritas Cluster Server (VCS)** – Eliminates planned and unplanned downtime by clustering critical applications and resources required by Informatica. VCS monitors and detects failures in PowerCenter and B2BDX resources with the Intelligent Monitoring Framework (IMF) without the overhead of a poll-based system. Agent-based monitoring and automation ensures clean application shut down and restart after a failure.
- **Virtual Business Services (VBS)** – Allows you to manage high availability and disaster recovery for integrated PowerCenter and B2BDX environments. VBS manages dependencies between the two applications and the associated repository database as well as the order in which the applications and their components are brought online in a start operation and taken offline in a stop operation. VBS is aware of the overall business service provided by the integrated PowerCenter and B2BDX solution and can take the appropriate action in the event of a failure to restore the entire service. This means faster recovery and minimal downtime, with no manual intervention.
- **Veritas InfoScale Operations Manager (VIOM)** – VIOM is a platform- and vendor-agnostic centralized management console for Veritas InfoScale Enterprise that also provides some visibility into other third-party infrastructure. Veritas InfoScale Operations Manager is used for monitoring, visualization, and management of system and storage resources. VIOM is also a reporting engine and can generate multiple reports, including a risk analysis report that can summarize issues that may arise within an environment that could reduce high availability and disaster recovery readiness of Informatica.

Figure 1 shows an overview of an Informatica PowerCenter HA and B2BDX HA integrated deployment managed by InfoScale Enterprise. InfoScale ensures that the PowerCenter and B2BDX resources and services are highly available in an active-active configuration. The Cluster File System leveraging Flexible Storage Sharing is mounted simultaneously on each node in the HA configuration to guarantee highly available clustered access to data.

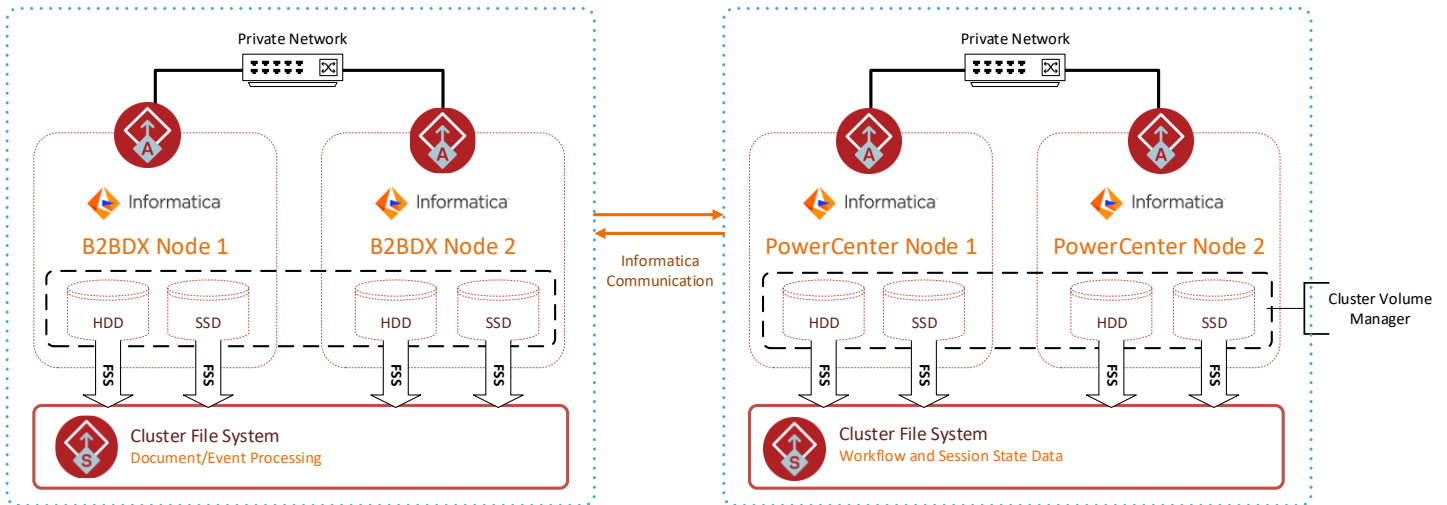


Figure 1 InfoScale Enterprise with Informatica PowerCenter HA and B2BDX HA

Availability

InfoScale Enterprise, combined with Informatica PowerCenter and B2BDX, enables a resilient, highly available and scalable EDI platform with the ability to maintain availability and data integrity across different data centers, geographical locations, or cloud platforms. The availability functions of InfoScale Enterprise enhance Informatica native capability to detect faults and automatically restore redundancy when a failover occurs. InfoScale Enterprise optimizes fault monitoring by consuming minimal compute resources with an Intelligent Monitoring Framework, ensuring that PowerCenter and B2BDX can be recovered almost instantly in the event of a failure with no impact on performance.

Maximum Availability for PowerCenter and B2BDX

InfoScale eliminates planned and unplanned downtime by clustering critical applications and the resources they require. InfoScale can monitor and centrally manage all the critical components of PowerCenter and B2BDX as well as underlying databases such as Oracle and MS SQL Server, to ensure maximum availability. By monitoring the status of PowerCenter and B2BDX resources and services and automatically restarting services in the event of a fault, InfoScale can dramatically increase their availability as well as that of any underlying databases. In addition to monitoring the application and database, InfoScale can instantly detect faults in the operating system, network, and storage resources.

InfoScale manages and automates all the operations required to detect faults, fail services over to standby nodes, and ensures that they can be failed back once the fault is resolved.

A critical component of InfoScale is the Agent for Informatica PowerCenter and B2BDX. This agent, developed specifically for Informatica, provides high availability for PowerCenter and B2BDX in either a standard or HA configuration. The agent can do the following to maximize application availability:

- In the event of a logical node failure, the InfoScale agent will automatically restart it. This increases the availability of the Informatica HA components which are responsible for managing all subordinate relative processes. For example, when using the PowerCenter component of the Informatica InfoScale Agent, it will start, stop and monitor only PowerCenter Service Manager which in turn holds responsibility for repository services,

web services hub, integration services, reporting services, and other application processes. If any Informatica logical node goes down for any reason, the InfoScale agent will automatically attempt to restart it.

- After a hardware failure has been resolved and the hardware is powered on, the InfoScale agent will automatically startup the Informatica service(s). The top level Informatica services will then re-join the Informatica services in use from the surviving nodes in proper membership order.

Regardless of whether there's a logical node failure or a hardware failure, Informatica HA (PowerCenter or B2BDX) will have already failed over master node responsibility to the next available Gateway host in the same domain. Informatica alone doesn't attempt to restart its applications on a failed node automatically. Therefore, the combination of InfoScale + Informatica HA yields the maximum possible application availability. This functionality behaves similarly with the B2BDX application components (dxserver, dxconsole, mft, and activemq). Figure 2 shows the Informatica components managed by the InfoScale agent for Informatica.

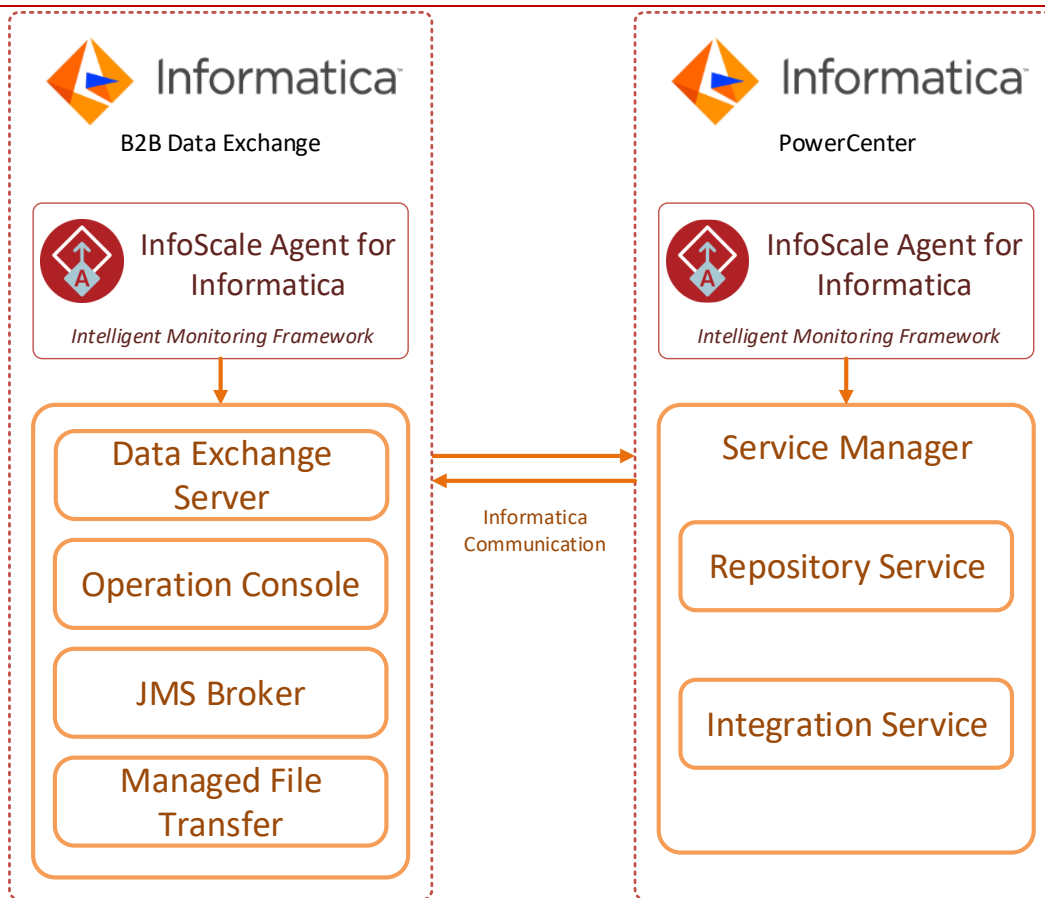


Figure 2 Veritas InfoScale Agent for PowerCenter and B2BDX

In addition to providing maximum availability for the PowerCenter and B2BDX, there is some other important functionality provided by the InfoScale agent for Informatica:

- The agent allows for priority to be established in an InfoScale cluster so that services may be shut down and started based on priority without manual intervention. This is extremely helpful during data center outages where

hosts in a PowerCenter domain are completely down. As soon as the hosts are powered on, the entire domain can be automatically recovered – assuming the database was also made highly available.

- InfoScale's **Virtual Business Services** feature allows for dependencies to be setup across InfoScale clusters so they can be managed as a single logical entity. This can be extremely helpful as it maximizes application availability by reducing the need for manual intervention for customers who typically setup B2BDX HA in one InfoScale cluster (with CFS for the shared file system) and the data exchange repository (relational database) in another InfoScale cluster. In this relatively common configuration, a VBS dependency can be established between each InfoScale cluster so that the Informatica service group (containing the Informatica resources) can be linked dependent on the data exchange database in a different InfoScale cluster
- After these linked dependencies are established, system administrators can be notified by InfoScale in the event of service group failures

Figure 5 shows an example configuration of an InfoScale Virtual Business Service for PowerCenter HA and B2BDX HA with the associated B2BDX repository database.

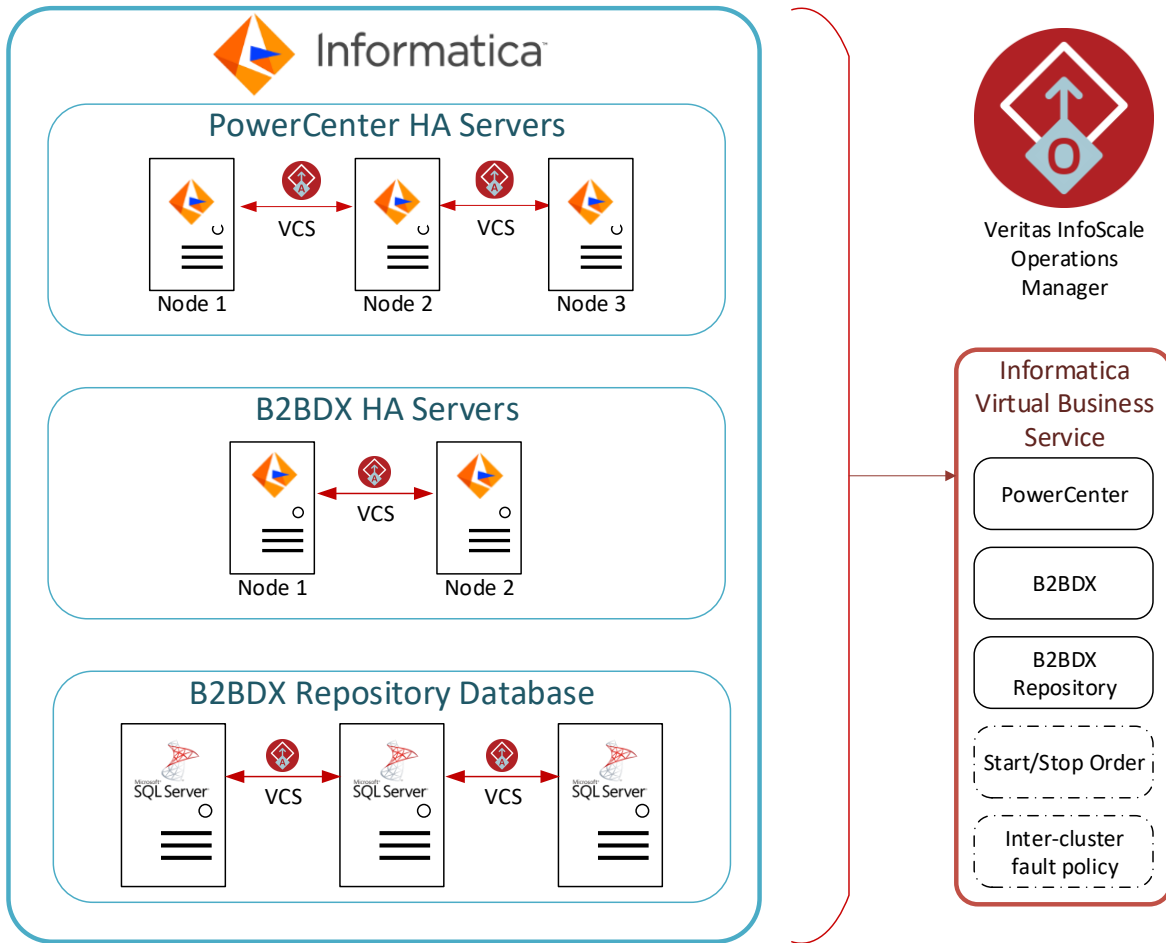


Figure 3 InfoScale VBS for Informatica

Intelligent Monitoring

The InfoScale agent for Informatica helps eliminate downtime and the potential for data corruption by using InfoScale's Intelligent Monitoring Framework (IMF), which reduces the time to action by allowing the agent to act immediately in the event of a system fault. The IMF for Informatica shown in Figure 4 allows the agent to register the process ID's of the Informatica processes with a kernel-level notification module that enables immediate (event-based) notification of resource state changes without having to periodically poll the resources to find the current state.

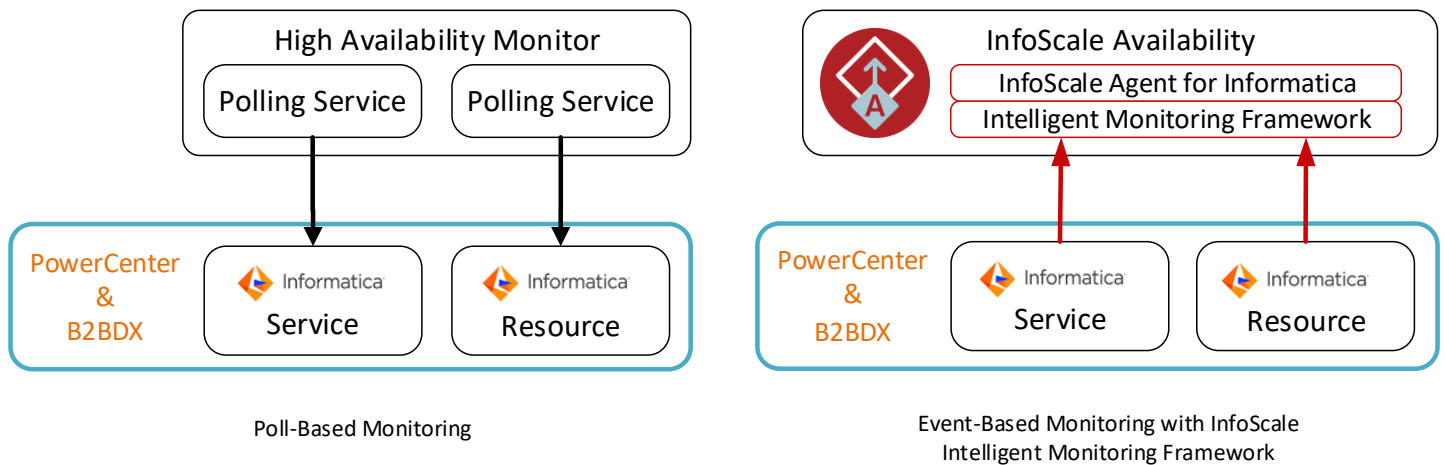


Figure 4 InfoScale Intelligent Monitoring Framework for Informatica

Performance and Optimization

While there are many storage options available, InfoScale Enterprise is recommended by Informatica as the storage solution of choice for [PowerCenter HA](#) and [B2BDX HA](#) because of its enterprise features, high performance, resiliency and scalability. The Flexible Storage Sharing feature of InfoScale Enterprise allows you to create a clustered storage environment using Direct Attached Storage (DAS) that can significantly improve application performance while reducing storage costs when compared with traditional Storage Area Network (SAN) environments.

Cluster File System (CFS)

The CFS component of InfoScale provides a highly available, fully POSIX compliant shared file system that delivers the availability and performance needed by an EDI platform. Using CFS in conjunction with the InfoScale agent for Informatica results in a highly available solution for PowerCenter and B2BDX with best in class performance and data resiliency – on premises or in the cloud.

The B2B Data Exchange components are configured on CFS for the best performance and reliability when processing documents and events. PowerCenter uses CFS to store data from the PowerCenter Integration Service that needs to be accessible to all nodes in an HA configuration that run the PowerCenter Integration Service. Figure 5 shows a two-node B2BDX HA configuration using Cluster File System provisioned on top of volumes managed by Veritas Cluster Volume Manager using disks shared with Flexible Storage Sharing.

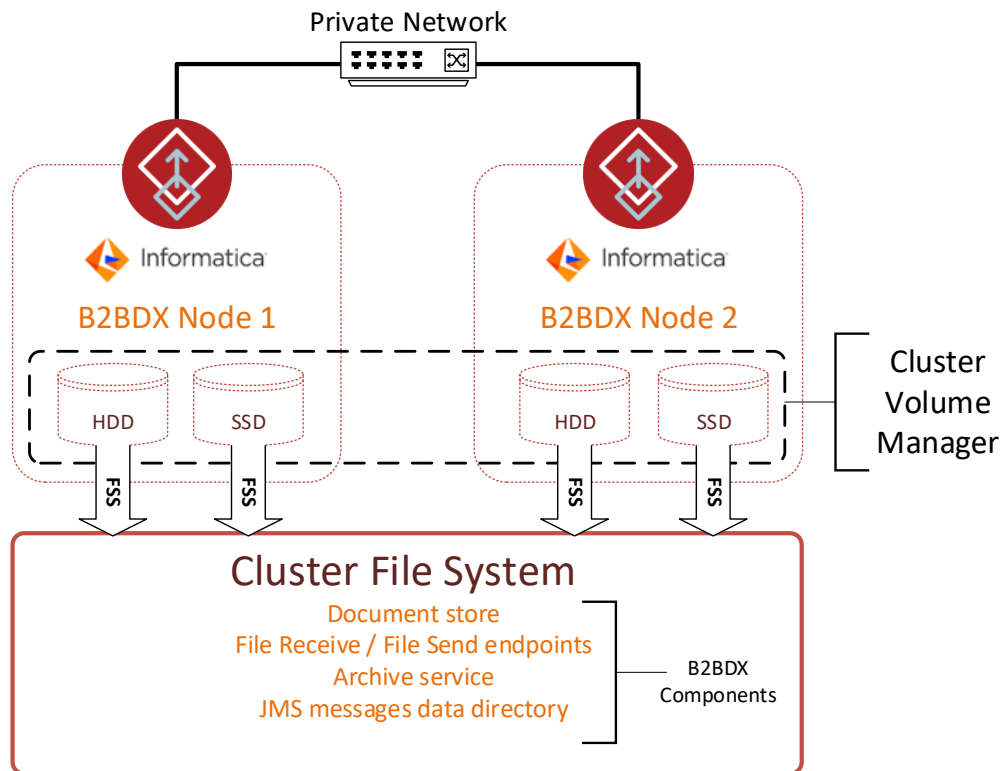


Figure 5 Cluster File System with FSS for B2BDX HA Configuration

NAS hardware with NFS as its file system is particularly difficult to implement in a solution that meets the criteria required for fault-tolerant shared storage. Using InfoScale CFS instead provides a simple solution that meets this requirement, providing fault-tolerant shared storage with write-order fidelity, synchronous write persistence, and distributed locking needed for Informatica HA configurations.

CFS enables you to mount the same file system on multiple Informatica nodes concurrently and is an extension of the industry-standard Veritas File System (vxfs). CFS uses a distributed locking mechanism called Global Lock Manager (GLM) to ensure all nodes have a consistent view of the file system. GLM provides metadata and cache coherency across multiple nodes by coordinating access to file system metadata, such as inodes and free lists. CFS advanced file lock management makes failover highly reliable.

Flexible Storage Sharing (FSS)

FSS is a feature of CFS that enables network sharing of local storage, cluster wide. The local storage can be in the form of Direct Attached Storage (DAS) or internal disk drives. Network shared storage is enabled by using a network interconnect between the nodes of a cluster.

FSS allows network shared storage to co-exist with physically shared storage, and logical volumes can be created using both types of storage creating a common storage namespace. Logical volumes using network shared storage provide data redundancy, high availability, and disaster recovery capabilities, without requiring physically shared storage, transparently to file systems and applications.

FSS in the Cloud

In public cloud environments where the options for shared disk devices are limited, FSS is the perfect solution for providing a clustered file system for Informatica PowerCenter and B2BDX deployments in the cloud. Public cloud storage devices are easily configured and shared between nodes by InfoScale Enterprise and make use of the same tools and interfaces as an on-premises, bare-metal installation.

Figure 6 shows an example of how CFS and FSS can be used for Informatica PowerCenter HA and B2BDX HA in public cloud environments.

I/O Fencing

A key requirement for protecting data in a cluster environment is to guarantee that there is always a single consistent view of cluster membership. In other words, when one or more nodes in the cluster stop sending heartbeats, the InfoScale software must determine which nodes can continue to participate in the cluster membership and how to handle the other nodes.

InfoScale I/O fencing has been designed to solve a common problem with any clustered storage solution, known as the “split-brain” condition, which occurs when there is a communication disruption between cluster nodes. InfoScale I/O fencing ensures that there is no data corruption in a split-brain scenario where cluster nodes cannot distinguish between a system failure and an interconnect failure, or where a node within the cluster is so busy that it appears to be hung and pauses communication with the other cluster nodes. To mitigate and resolve the split-brain condition, Veritas InfoScale’s I/O Fencing system guarantees the data integrity of the B2BDX and PowerCenter data stored on the InfoScale CFS by determining which nodes in the cluster should remain in the event of a communication disruption. When a disruption occurs, the node which has failed is ejected from the cluster and prevented from accessing the data disks.

InfoScale I/O Fencing can be configured to work on-premises, on various hypervisors, or in public cloud platforms. Informatica’s recommended approach is to configure CFS to use InfoScale I/O fencing for all the nodes in the cluster.

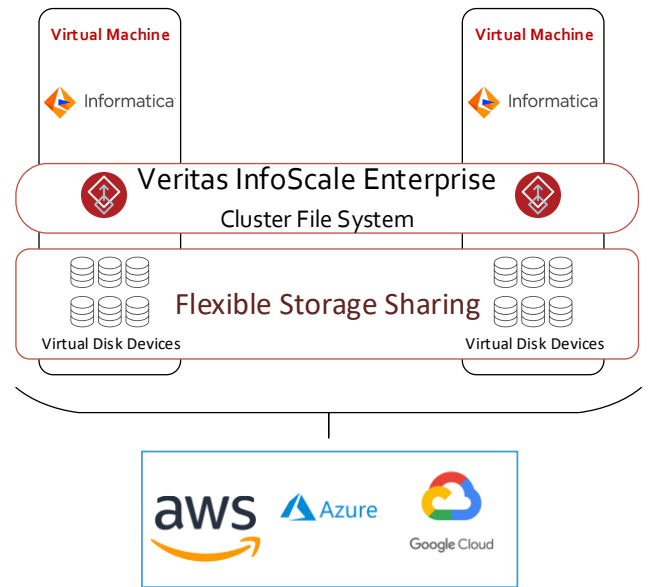


Figure 6 Flexible Storage Sharing in public cloud

Efficiency and Agility

InfoScale simplifies the administration required for high availability and shared storage in Informatica environments with a single interface and a single suite of tools regardless of OS or platform, reducing complexity and increasing operational efficiency. InfoScale Enterprise automates the process of conducting a failover by assessing the state of the application and storage, performing the appropriate actions required to fail over PowerCenter and B2BDX nodes to standby nodes, and bring up the necessary services to ensure a successful failover.

Veritas Volume Replication and Global Cluster Option

Veritas Volume Replicator (VVR) provides organizations with an extremely flexible storage hardware independent alternative to traditional array-based replication solutions as well as a robust mechanism for moving data into public cloud environments. VVR delivers the flexibility of block-based continuous replication as well as file-based periodic replication with Veritas File Replicator (VFR).

VVR also manages and maintains write-order fidelity, which is an important feature that ensures data consistency and protects against data corruption by ensuring that the data on the secondary volume is consistent with the data on the primary volume. This guarantees that the application is working with the same data at both sites. Native storage replication solutions do not typically maintain write-order fidelity. VVR maintains write-order fidelity by logging write operations on the primary volume in the order in which they are received and applying them on the secondary volume in the same order.

InfoScale facilitates disaster recovery for PowerCenter, B2BDX and the associated repository database between geographically dispersed sites by integrating VVR replication with the Global Cluster Option (GCO). GCO extends InfoScale from local cluster failover to inter-cluster failover between different data centers or public clouds (see Figure 8 InfoScale Public Cloud Support).

Veritas InfoScale Operations Manager

Veritas InfoScale Operations Manager (VIOM) is a platform and vendor agnostic centralized management console for InfoScale and third-party infrastructures. Veritas InfoScale Operations Manager is used for monitoring, visualization, and management of system and storage resources on-premises or in the cloud. VIOM is also a reporting engine and can generate multiple reports, including a risk analysis report that can summarize issues that may arise within an environment that could reduce high availability and disaster recovery readiness.

A typical VIOM deployment consists of 2 main components: a management server and managed hosts. Depending on the usage scenario, VIOM may also discover virtualization environments, SAN/NAS infrastructure as well as SAN fabrics. Managed hosts can be added into VIOM either using agents or as agentless hosts in the situation where there is no InfoScale software installed on the target hosts. VIOM manages agentless hosts using SSH or WMI. The level of visibility within VIOM for agentless hosts is the infrastructure only; there is no application visibility for agentless hosts.

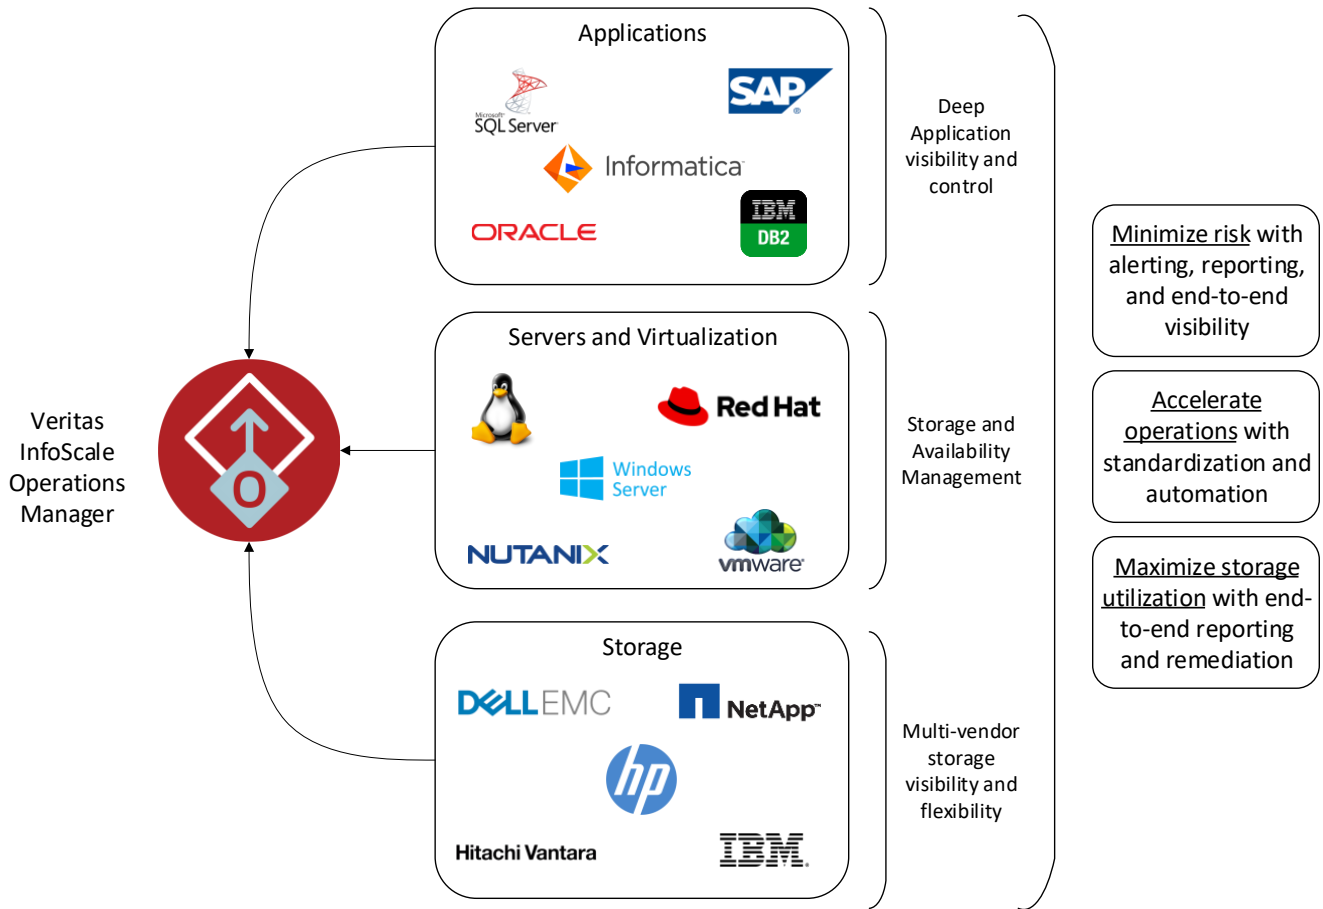


Figure 7 Veritas InfoScale Operations Manager Integration Points

Public Cloud Support

The Veritas InfoScale Enterprise product suite helps make Informatica PowerCenter and B2BDX highly available and resilient both on-premises and in public cloud platforms like Amazon Web Service, Microsoft Azure, and Google Cloud Platform. InfoScale protects PowerCenter and B2BDX systems from unforeseen failure at both the application layer down to the infrastructure layer with tight integration into public cloud platforms, all with a similar look-and-feel to reduce operational complexity.

In addition to protecting your data and enhancing performance in a single cloud provider, Veritas InfoScale provides the option of protecting your data between public clouds and on-premises data centers. InfoScale's ubiquitous OS and platform support allows you




Figure 8 InfoScale Public Cloud Support

to move data seamlessly and reliably while ensuring that all InfoScale functionality is available when you need it.

In case of a failure within on-premises resources or a public cloud, InfoScale enables you to maintain seamless support for failing over to another public cloud (or on-premises/nearline datacenter).

Figure 8 shows Informatica installations in multiple public clouds as well as on-premises. InfoScale Enterprise supports high availability topologies where applications such as PowerCenter and B2BDX can fail over between cloud and on-premises environments.

Summary

Informatica PowerCenter and B2BDX are leading solutions for data integration and transformation that enable enterprise data warehousing and flexible system-to-system communications across different technologies. Veritas InfoScale Enterprise offers a proven end-to-end, integrated solution for ensuring high availability and robust shared storage for PowerCenter and B2BDX deployments. The combination of InfoScale Enterprise's industry-leading features with Informatica provides several key benefits:

- ✓ Maximum uptime for Informatica enterprise data integration and transformation initiatives
- ✓ High performance and resilient clustered storage environment for PowerCenter and B2BDX HA configurations
- ✓ Simplified availability management for multi-tiered Informatica EDI services with Virtual Business Services

As the solution recommended by Informatica for storage and high availability with PowerCenter HA and B2BDX HA, InfoScale Enterprise is a proven platform that enables the best possible availability and uptime for your business critical Informatica EDI projects.

Table of Figures

Figure 1 InfoScale Enterprise with Informatica PowerCenter HA and B2BDX HA.....	5
Figure 2 Veritas InfoScale Agent for PowerCenter and B2BDX	6
Figure 3 InfoScale VBS for Informatica	8
Figure 4 InfoScale Intelligent Monitoring Framework for Informatica	9
Figure 5 Cluster File System with FSS for B2BDX HA Configuration	10
Figure 6 Flexible Storage Sharing in public cloud.....	11
Figure 7 Veritas InfoScale Operations Manager Integration Points	13
Figure 8 InfoScale Public Cloud Support	13

ABOUT VERITAS

Veritas Technologies is a global leader in data protection and availability. Over 50,000 enterprises—including 99 of the Fortune 100—rely on us to abstract IT complexity and simplify data management. Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas supports more than 500 data sources and over 150 storage targets, including 60 clouds. Learn more at www.veritas.com. Follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

2625 Augustine Drive, Santa Clara, CA 95054
+1 (866) 837 4827
veritas.com

For specific country offices
and contact numbers,
please visit our website.

