

VMware Delivers a Modern Approach to Data Management

High-Speed Data, Elastic Scale and Database-as-a-Service Capabilities for the Cloud Era

KEY HIGHLIGHTS

- VMware vFabric™ GemFire® provides high-performance data management with in-memory speed and elastic scalability.
- VMware vFabric SQLFire is designed to handle structured data and is well-suited for IT teams with relational database experience using a familiar SQL interface.
- VMware vFabric Data Director and VMware vFabric Postgres enable a database self-service solution that leverages virtualization to increase IT agility and improve enterprise governance.
- vFabric solutions are designed for IT organizations to confidently manage data across private, public and hybrid clouds.

Extract More Value While Managing the Growth of Enterprise Data

Application modernization transforms business. Today's modern applications—characterized as Web-oriented, data-intensive, dynamic, virtualized and real-time—both consume and generate data. Yet, existing approaches to data management prevent many businesses from harnessing the true value of their data.

To meet increasing demands for immediate information, enterprise IT organizations realize they must effectively address high growth of their data volumes, poor scalability and performance of their monolithic databases, and data-access latency. Even as they upgrade performance and scale, IT teams are tasked with reducing the high costs of traditional relational database management systems (RDBMSs), while improving distributed data access, simplifying database provisioning and governance, and deploying cloud data strategies and virtualization.

At the same time, line-of-business owners want to leverage the untapped value in enterprise data to create competitive advantage for the business. They seek to generate incremental revenue by minimizing latency during online customer experiences. They request situational data and alerts so they can improve decision making and risk assessment, as well as respond immediately to opportunities and threats. They want IT teams to ultimately deliver the single view of complex business transactions that their businesses demand.

The exponential growth in data to support today's applications poses challenges for IT and business leaders. It has led to unprecedented database sprawl—the proliferation of undermanaged, undersecured and even unknown databases across organizations. With manual operations for this growing set of heterogeneous databases restricting developers' ability to rapidly access and provision data services, organizations are looking for a better provisioning approach—one that delivers self service while keeping IT in control.

Successful application modernization begins with the right approach to data management. VMware®, the global leader in virtualization and cloud infrastructure, delivers a modern approach to data management that enables businesses to make accurate decisions based on real-time, event-driven information. At the same time, IT organizations gain the benefits of elastic scale and database-as-a-service (DBaaS) capabilities. For both business and IT leaders, the VMware approach reduces costs compared to traditional database architectures.

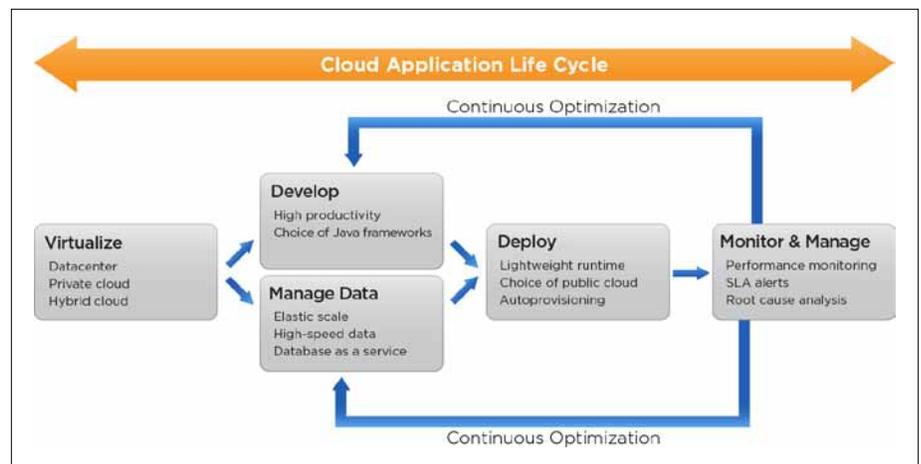


Figure 1. Enterprise data management is an integral part of application-modernization initiatives and plays a strategic role in the cloud application life cycle.

Distributed caching platforms (DCPs) have proved to be key enablers for mainframe and DBMS offloading, performance and scalability boosting in Web commerce applications, high-performance publish and subscribe, cloud and software as a service (SaaS), and other scenarios.

— Fabrizio Biscotti and Massimo Pezzini
*Emerging Technology Analysis:
Distributed Caching Platforms
Enable Innovative Scenarios*
Gartner, Inc.
November 2, 2010

Go Beyond Existing Database Architectures

IT organizations and enterprise stakeholders want modern business applications that run fast and process large data volumes. However, most enterprise data is locked in legacy mainframes and traditional RDBMSs that were designed decades ago. IT leaders and line-of-business owners require a modern approach to data management that keeps up with the speed of business by enabling the following capabilities:

- **Dynamic scalability without higher costs** – The amount of data flowing through enterprises is predicted to grow by 50–60 percent a year.¹ Traditional RDBMSs can only scale vertically, in very large increments. As a result, IT teams can expand services only by adding bigger machines and purchasing more licenses, which increases cost and management overhead. IT organizations want to scale quickly, incrementally and dynamically (both up and down) in response to variable data-consumption demands (e.g., seasonal peak usage).
- **Real-time data and event-driven communication** – Whereas large data warehouses are ideal for analyzing historical data, many modern applications must process up-to-the-second information to detect business risk or threats and missed revenue opportunities. Today's enterprises require rapid, real-time information dissemination—activated by event detection and alerting—to achieve superior decision making and effectively assess risks.
- **Zero Web site latency** – When customers experience downtime or slow online service, business suffers. Traditional RDBMSs require a round trip to the server, which increases data latency. Business owners demand better performance to generate incremental revenue through the enhanced customer satisfaction that results from minimizing Web site latencies.
- **An accurate view of all transactions** – Business is complex, but decision makers want one consistent view of the dynamic data distributed throughout the enterprise, across regions and time zones. To provide this view, IT organizations need a better way to synchronize distributed data sources so that updates that occur at one end point immediately propagate to all other end points.
- **Faster and more affordable mainframe processing** – Traditional mainframe applications typically run batch processes, which can take several hours to complete and slow business transactions. An even more pressing issue for today's enterprises is the cost of high-volume mainframe queries. Even a midsize company can spend hundreds of millions of dollars each year on accessing mainframe data. As enterprises set application-modernization goals to reduce costs by mitigating the need to access mainframe data, they want to leverage modern operational data stores that are both elastic and real-time.
- **Simplified database provisioning and governance** – Databases are time-consuming to provision. Without automated policies and role-based access, they also make it difficult to maintain good corporate governance and eliminate database sprawl. By taking advantage of virtualization and cloud computing capabilities, IT organizations aim to provide self-service database provisioning to their developers and enable DBaaS.

¹ IDC. "Worldwide Data Protection and Recovery Software 2010–2014 Forecast: Cloud, Deduplication, and Virtualization Stabilize Market" Robert Amatruda. August 11, 2010.

EXAMPLES OF A MODERN DATA-MANAGEMENT APPROACH

Travel – A global online travel agency dramatically cut Web site response time for its users, improving its look-to-book ratio, which increased customer satisfaction that yielded millions in additional revenue each year.

Securities – Traders are leveraging a high-speed data backbone to enable a single, consistent view of transactions across global exchanges and to support their real-time pricing and risk models so they can react immediately to market fluctuations.

Healthcare – A nationwide pharmacy benefits network is reducing the cost of its mainframe applications by deploying an in-memory database as its system of record while providing a real-time view of patient claims.

Financial services – One of the largest global financial-services firms has achieved elastic scale across tens of thousands of nodes in perfect sync with peaks and valleys of demand.

Government – The U.S. Defense Information Systems Agency (DISA) maintains real-time situational awareness of military operations worldwide with more than 60,000 data updates each minute.

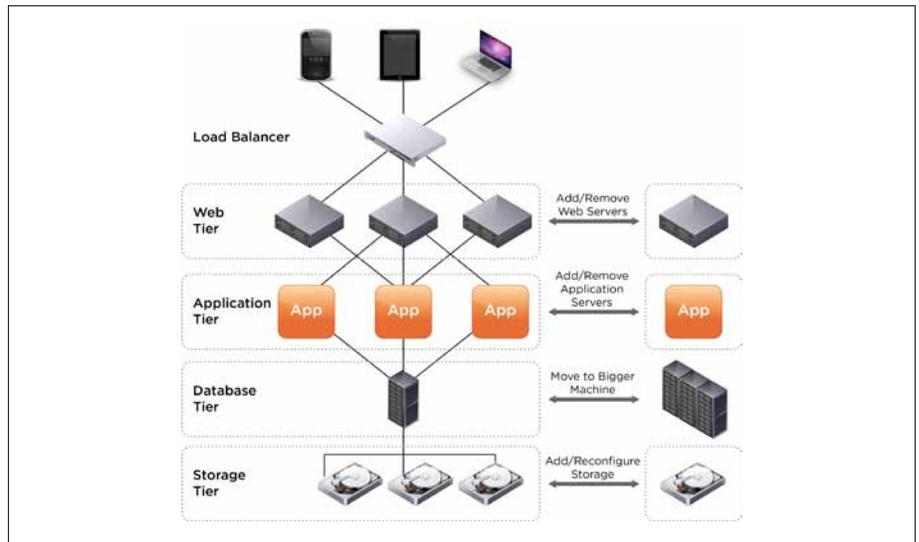


Figure 2. Traditional approaches to scaling the database tier require IT teams to move to larger machines and purchase more licenses, which increases IT cost and management overhead.

Build on a Modern Data-Management Foundation

VMware vFabric GemFire is the foundation for a modern approach to data management that meets IT and business leaders’ needs. A distributed data-management platform, GemFire eliminates data-access latencies and provides dynamic scale in support of modern enterprise and cloud applications. It delivers the reliability of a database with the performance that only in-memory data architecture can provide.

GemFire can be used as an operational data store and system of record. It can complement existing databases by accelerating queries and response times or serve as a lower-cost, highly scalable alternative to the monolithic RDBMS. GemFire lowers TCO—by leveraging commoditized x86 hardware and minimizing RDBMS license costs—while providing superior reliability and quality of service. For example, while it ensures fault tolerance, high availability and transactional integrity, GemFire can optionally write or overflow data to disk for backup and recovery. It does so by using an innovative memory-disk hybrid approach. With GemFire, enterprises can also leverage virtual infrastructure, as well as other data stores deployed in private, public or hybrid cloud environments.

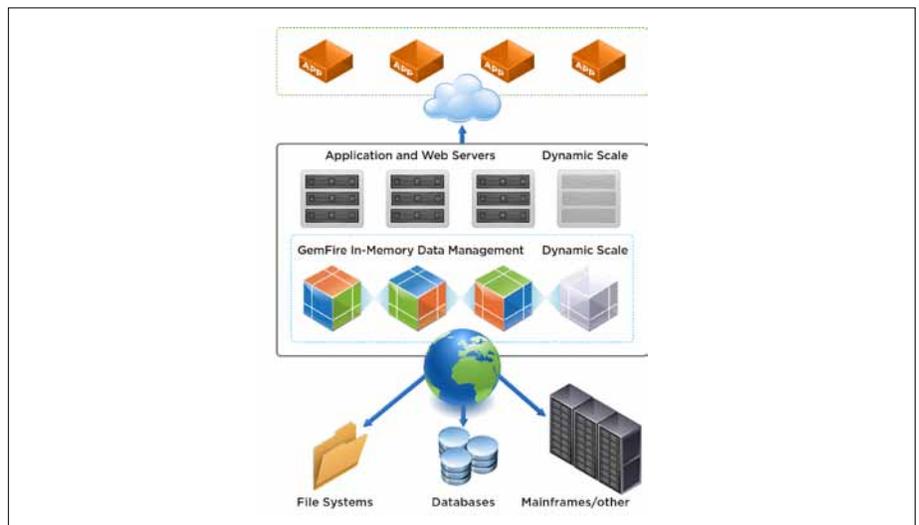


Figure 3. To overcome monolithic RDBMS scale, latency and cost challenges, GemFire delivers both speed and reliability using an innovative memory-disk hybrid approach.

The use of data fabrics provides the opportunity to build newer, decision-centric applications. These types of applications consume large volumes of data at high speeds to perform the analysis needed to detect when a decision needs to be made and also perform analysis to determine best options for the decision.

— Al Hilwa and Maureen Fleming
Modernizing Application Platforms for Cost-Effective Cloud Readiness
IDC
June 2011

Today's enterprises are embracing the GemFire in-memory data architecture for global data grids and cloud data management because it provides the following critical capabilities:

- **Delivers high-speed data processing** – By managing data in-memory, GemFire can perform database computations without the constraints imposed by database disk I/O performance. Its in-memory design drives high-performance capabilities so that ultra-high-volume data updates can be processed each second and a response can be delivered to users in real-time, improving the accuracy of decision making.
- **Enables elastic scalability** – GemFire provides dynamic scalability while lowering costs. By adding or removing x86 machines on the fly—in just minutes without more licenses—GemFire enables IT teams to easily and quickly scale systems up to meet peak or seasonal demand, and back down when demand subsides.
- **Provides an enterprise or global data backbone** – In a connected world, the enterprise-wide or global distribution of dynamic data is critical to business and government operations. Using GemFire and WAN gateways, IT organizations can synchronize distributed systems located across regional datacenters or continents, and manage 24x7 global workflows.
- **Achieves a single, consistent view of complex transactions** – With GemFire, changes to data at one location are immediately reflected in all other locations where the data element is used. This instant updating provides decision makers with a complete, accurate view of transactions that rely on dynamic data from multiple sources.
- **Enables real-time, situational awareness and event-driven detection** – GemFire can monitor sudden changes in globally distributed data to detect certain events—for example, significant changes in stock prices or troop movement in foreign lands. It can alert and update appropriate personnel or initiate automated processes and responses.
- **Reduces the cost of mainframe applications** – GemFire mitigates costly dependencies and bottlenecks by offloading data from mainframe responsibility for both online and batch processing. As an enterprise processes and stores data in GemFire, requests bypass the mainframe. This path dramatically lowers mainframe application-processing costs and reduces batch-processing times from hours to a minute.
- **Offers an operational data store with database reliability** – GemFire can be configured as a distributed database—running over large and small clusters or other grids—while providing the persistence layer for in-memory operations. With GemFire, IT teams gain not only high performance, but also database-like persistence to ensure that disk failure on one node does not result in data loss.

Enable Fast, Scalable SQL Data in the Cloud

VMware vFabric SQLFire is an in-memory SQL database designed to address many of the same challenges as GemFire. Whereas Java developers appreciate GemFire's ease of use because of its tight integration with the Spring development framework, the SQLFire configuration and deployment model is ideally suited for anyone with relational database experience. IT organizations can now easily leverage existing database administrator (DBA) skill sets and tools when managing a variety of databases deployed internally or in the cloud.

Like GemFire, SQLFire can be used either as the primary data store or as a front-end data-management layer for one or more existing databases. It works effectively with a large ecosystem of compatible products and frameworks, including object-relational mapping tools (e.g., Hibernate and NHibernate), schema-editing and database management tools (e.g., SquirrelSQL), database-replication products, Spring Java Database Connectivity (JDBC), and more.

Enterprises must support hundreds or even thousands of applications to meet growing business demands, but this growth is dramatically driving up the cost of running and managing the databases under those applications. The stress this puts on the IT budget makes it harder to provide databases to support new requirements such as Web 2.0 applications or other emerging collaboration solutions or even to support more mundane uses such as increased application testing. A new emerging option called database-as-a-service hosts databases in the cloud and is a good fit for some new apps.

— Noel Yuhanna
Database-As-A-Service Explodes on the Scene
Forrester Research
July 31, 2008

Address Database Operations, Governance and Virtualization

Modern applications both consume and generate data, requiring enterprises to manage more information than ever. This has led to unprecedented database sprawl. Although IT organizations want to reduce the burden and improve governance, inflexible enterprise database operations—including manual provisioning and deprovisioning, cloning and backup, and restoring databases—are time-consuming and difficult-to-manage processes that hinder business agility.

For example, IT teams can take several days or weeks to provision a new database for a development or quality assurance (QA) test environment, yet that database might be used only for several days of testing. This inefficient provisioning process is frustrating to both development teams and DBAs. So instead of waiting for internal IT organizations to provision new databases, developers, QA organizations and other business stakeholders are now deploying their own databases or turning to public cloud providers to meet their needs.

In today's organizations, many x86 servers run databases. However, most of them are undermanaged and unprotected by high availability or backup solutions because they were not provisioned by internal IT teams. Without internal IT oversight, these databases affect productivity and consume costly resources.

With server computing power continuing to grow, many organizations are moving toward database consolidation. Virtualization, as the foundation for cloud computing, offers an easy and efficient solution for database consolidation that improves quality of service, accelerates the database life cycle from development to production, and reduces infrastructure and software license costs.

Power DBaaS for the Cloud

With database consolidation through virtualization, enterprises only begin to achieve the cost and operational benefits of cloud computing. This is primarily because at their core, monolithic databases still operate as if they were running on dedicated physical machines with unchanging resources and because consolidation requires process changes, including automated provisioning. Optimized for VMware vSphere®, the industry-leading virtualization platform, VMware vFabric Data Director and VMware vFabric Postgres deliver the IT agility and enterprise governance required for organizations to achieve the full benefits of virtualization and cloud computing.

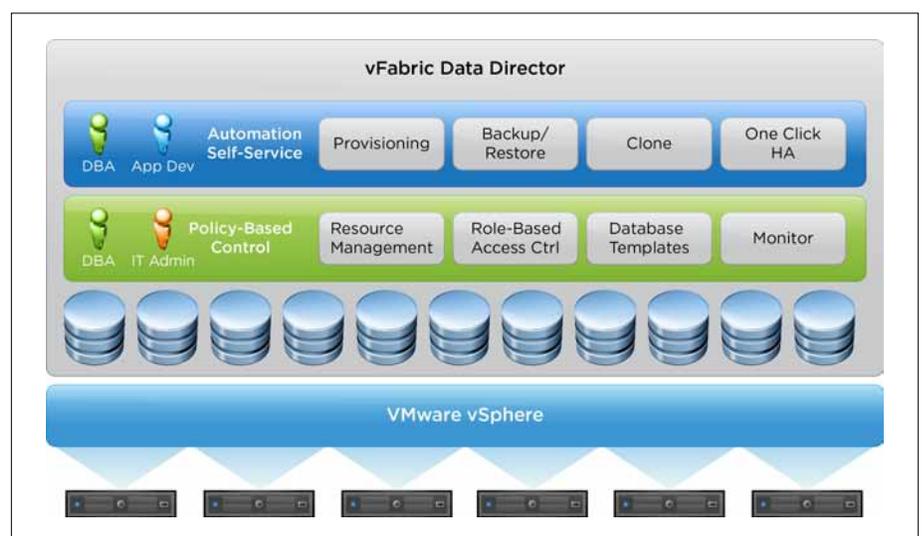


Figure 4. VMware data-management solutions enable your IT organization to power DBaaS for your cloud.

vFABRIC SOLUTIONS

In addition to GemFire, SQLFire, Data Director and Postgres, the core services in the vFabric Cloud Application Platform include:

Spring – The most popular Java development framework in the world

vFabric tc Server – An enterprise version of Apache Tomcat optimized for Spring and vSphere

vFabric RabbitMQ™ – The leading open-source implementation of the Advanced Message Queuing Protocol, enabling a cloud-ready approach to messaging

vFabric Hyperic® – Proactive performance management of custom applications through transparent visibility into modern applications deployed across physical, virtual and cloud environments

Cloud Foundry – The industry's first open platform as a service (PaaS)

In contrast to current database provisioning and configuration tools and processes—which take can take weeks and slow innovation cycles—vFabric Data Director speeds development cycles while keeping IT in control. It offers just-in-time provisioning with an easy Web interface and templates to guarantee proper configurations, so IT teams can provision fully configured databases in minutes. vFabric Data Director also includes built-in, automated backup policies to enable, monitor and test backup and recovery policies for all databases, saving DBAs time and enhancing business agility. It increases developer productivity through innovative linked database clones. Rather than waiting days for clones to be created, developers can now use a self-service portal to get a clone up and ready in minutes regardless of database size, dramatically increasing developer productivity.

Data Director supports fine-grained resource management, as well as custom and out-of-the-box role-based access control. With these features, IT organizations can enforce policy throughout a secure environment. More than 30 different security privileges enable self service and ease security management. The flexible and convenient database templates included with vFabric Data Director are simple to use, and they ensure provisioning and configuration reliability and repeatability without requiring sophisticated DBAs, who have limited time.

Data Director extends all of the benefits of virtualization to databases through database-aware high availability and automatic database configuration. Affordable and easy to configure and manage, Data Director offers zero-data-loss protection and automatic, administrator-free failover from host, operating-system and database failures.

Postgres is the first database enabled on the Data Director platform. Additional commercial and open-source databases will be enabled on the Data Director platform in future releases. Postgres is based on—and fully compatible with—PostgreSQL, an advanced open-source database used widely by enterprises today. ACID (atomicity, consistency, isolation, durability) and American National Standards Institute (ANSI) SQL compliant, Postgres is a relational database optimized to run in the vSphere environment. It ensures consistent database performance with VMware patent-pending elastic database memory technology. Together, Data Director and Postgres offer a cloud-ready infrastructure.

With Data Director and Postgres, IT teams can replace the plethora of undermanaged and unprotected databases in their organizations with a powerful self-service solution. Together they can power DBaaS for your cloud—the right cloud model to match your business needs—improving agility while maintaining IT control and compliance.

Learn More About Data and Application Modernization

Database modernization is a critical step in the application-modernization process. With VMware data-management solutions, your IT organization can modernize your data architecture and achieve greater IT and business agility—and you can confidently manage data across your cloud.

GemFire, SQLFire, Data Director and Postgres are core components of the VMware vFabric Cloud Application Platform—an integrated application development and delivery solution for virtual and cloud environments. Enterprises choose proven vFabric solutions to increase agility while reducing the complexity and operational costs of the entire application life cycle.

Learn more about how your organization can begin to modernize its business applications and take a modern approach to data management. Schedule a consultation with a VMware data-transformation expert or visit <http://www.vmware.com/data>.

