White Paper

Load Balancing Oracle Database Traffic

Databases are essentially the heart of the IT infrastructure that today's business runs on. Historically, databases have not had an efficient load balancing system; but technology has advanced enough now that organizations can load balance many facets of transactional database systems such as Oracle databases.

by Don MacVittie
Technical Marketing Manager

Contents

Introduction

There is very little debate about the importance of databases in the corporate data center—without them, business would grind to a halt. Unstructured data is growing at a much faster pace than structured data, but structured data represents an organization's accumulated knowledge about customers, orders, suppliers, and even employees. Yet effective load balancing for mainstream database management systems (DBMSs) has escaped the industry for many years. This is partially due to the transactional nature of DBMS traffic, and partially to the critical nature of databases. Anything that inserts another potential point of failure between databases and the applications they service has been viewed with a high level of skepticism.

Advances in database technology and the proven track record of Application Delivery Controllers (ADCs) have merged to change the face of the marketplace. Once database clusters became relatively common, it was a matter of time before users realized that clustering is, in many senses, software-implemented load balancing. In the meantime, ADCs came of age, offering load balancing and a whole host of other functionality from monitoring to security. The number of applications sitting behind ADCs, combined with the growth in database clustering and increasing desire for high availability solutions at the database level, naturally led to organizations using ADCs to balance the workload of DBMS products.

Some more cautious organizations utilize ADCs to speed access and switchover for DBMSs; other, less risk-averse organizations are pushing the boundaries with outright DBMS load balancing. Organizations with larger database workloads utilize database clustering, while those with smaller loads generally approach the problem from a stand-alone database perspective.

In any organization, IT staff must determine whether load balancing databases is in their best interests and if so, which features are best suited to their architecture. F5 products provide various options for load balancing these highly complex critical systems, so organizations can ensure their DBMS architectures are more secure, fast, and available.

F5 Speeds Oracle

"We have been able to deploy Oracle globally and mitigate the effects of latency due to distance with the web acceleration technologies implemented in F5 products. Oracle performs better and more predictability for our users throughout the world."

Senior IT Architect, Large Enterprise Construction Company

Source: TechValidate

TVID: D48-242-166

Database Management Systems

Database management systems rely on network connections to do their tasks in support of applications. This makes them a natural target for load balancing at the network level.

But there are significant challenges to load balancing DBMSs. First and foremost, a DBMS is assumed to have access to all of the records for a particular table, which implies that the database is updated directly. When load balancing across DBMSs, how can it be arranged such that all instances have access to all data for all tables? DBMSs also require transactional integrity to guarantee that all of the changes relevant to a transaction complete, or else the entire transaction doesn't complete. Transactional integrity has been one of the limiting factors of DBMS load balancing. If load is being distributed across multiple databases, how does IT guarantee that all of the elements of a single transaction go to a single instance of the database so that transactional integrity is insured?

When IT utilizes clustered databases, these issues are handled at the clustering software layer. The software ensures that that each instance has access to the entire database, and sends connections to the correct instance.

But there is always room for improvement, and clustering is no exception. When Oracle database clusters are deployed, a server that encounters problems and goes offline may take a significant amount of time to notify applications. Applications that are Oracle Fast Application Notification (FAN) enabled will be notified quickly, while other applications—the bulk of the application infrastructure—will take much longer to realize there is a problem and reconnect to the cluster to get access to a valid server.

Load Balancing Clustered Databases

Load balancing clustered databases isn't actually load balancing, per se, but rather a way to create a highly available infrastructure between database clusters. F5® BIG-IP® Local Traffic Manager™ (LTM), an ADC, uses a variety of monitors to check the health of pool members, so if the primary and secondary clusters are configured as members of a single pool and utilize priority queuing, when the primary goes down, the secondary will automatically receive the traffic. This is one small bit of a complex architecture, but it is an enabling part that automates failover so that there is no delay while administrators are notified of a problem, go to look

into the problem, and then manually make the exact same switch. Since BIG-IP LTM provides the connection address for applications utilizing the database (as a virtual IP in front of the pool), switchover doesn't require any IP address juggling exercises on either server or client applications.
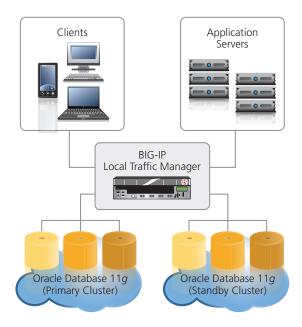

Figure 1: BIG-IP LTM manages failover for clustered Oracle Database 11*g*.

This is the easiest solution to implement because there are no heavy data or software architecture requirements beyond the choice to use high availability clusters. Using multiple clusters, without BIG-IP LTM, requires that IT have a replication system in place that is near real time, or the idea of failover won't work to begin with. There must be a mechanism for that replication to be two-way, so that whichever system is active is feeding back to the one that is not. All of these are requirements of utilizing multiple clusters, not of using BIG-IP LTM to provide a high-performance, highly tunable failover between the clusters.

Load Balancing All Databases

Without BIG-IP LTM, applications that conform to Oracle's FAN failover system can fail over quickly and gracefully. BIG-IP LTM extends that failover ability to all database applications. Given the number of applications that do not support FAN, this is a huge benefit in the short term. BIG-IP LTM achieves this with two automation tools. The first is a set of iControl scripts, which extend FAN to the

BIG-IP system by marking a node as down on the BIG-IP device if FAN reports it as
down, and up if FAN later reports it as being back up. The second automation tool
is built into BIG-IP LTM, and is an easy-to-use configuration setting that instructs the
BIG-IP device to reject connections to devices marked as down. Since BIG-IP LTM is
a full TCP proxy, if this configuration setting is turned on, when FAN marks a node
as down, it is reflected in the status of the node on BIG-IP LTM; thus connections
attempting to reach the downed node are rejected by the BIG-IP device. This starts
the process of the application reconnecting to a new database server that can
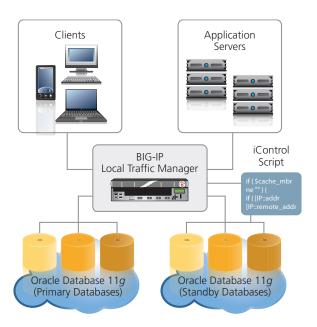handle application requests.

Figure 2: BIG-IP LTM extends FAN notifications to all applications, not just those built on the
Oracle JVM.

With BIG-IP LTM standing between the application and the databases, acting as
a full TCP proxy with knowledge of the state of database servers, connections can
be reset immediately upon attempting to communicate with a downed server. This
can happen when a server goes down in the middle of a communications stream.
BIG-IP LTM marks the database as down, and when the next request comes from
the server, BIG-IP LTM resets the connection, forcing the application to a different
database upon return. For applications that are not FAN-enabled, Oracle uses
industry-standard TCP timeouts as the notification mechanism. While this offers
the broadest possible support for applications, it is too slow for many environments,
as the application has to send a request and then wait for the TCP timeout interval
before determining that it must reset the connection.

BIG-IP LTM also offloads monitoring from Oracle. From BIG-IP LTM, a single copy of the SQL query utilized to check the status of Oracle databases can be applied to all Oracle instances. This reduces the opportunity for error by removing many redundant copies of this script from around the network. It also reduces network traffic and management time by enabling IT staff to control frequency or pings from a centralized location via health monitors built in to the BIG-IP system and the query designed to test Oracle status.

And as with all applications placed behind BIG-IP LTM, if administrators need to perform maintenance, connections to the database can be gracefully bled off of a single database server until there are zero connections. There is no need to kill off all active connections to take the server down; rather, the administrator can just mark it as not accepting new connections, and let the connections slowly drain away as each is completed. In the case of an Oracle Real Application Cluster (RAC), this would have the effect of sending new connections to the other servers in the cluster. In a standalone database environment, this would have the similar effect of shipping all connections to the redundant database(s). When maintenance is complete, the administrator can return the server to the pool, and it will resume accepting connections as if it had never left.

In a nutshell, BIG-IP LTM gives organizations faster connection resets when a database or entire cluster goes offline, centralized management of SQL scripts for testing, extension of FAN to non–FAN enabled applications, and the ability to take servers out of the pool to perform maintenance or even replace the hardware.

Replication Enhancement

It is impossible to load balance applications across databases unless those databases are synchronized in some manner. While there are a variety of ways to handle replicating the contents of a database, by far the most common is to make one database the master and one the secondary, then replicate changes to the master through to the secondary. This process is well supported by Oracle and third parties, and works with varying degrees of success depending on the situation. In general, as the distance the data has to be transported and the volume of that data both grow, the more performance of applications designed for replication degrades. Since most of the replication applications on the market today have their roots in LAN-only replication, this is not surprising; but replication over the WAN is becoming more prevalent, causing major problems.

Oracle offers many options for replicating databases, and these products work very well over the LAN. However, these same products perform less well over the WAN, where there are a whole different set of points at which performance can degrade. BIG-IP® WAN Optimization Manager™ (WOM) helps products like Oracle GoldenGate speed data replication from one data center to another by enhancing the performance of the WAN. In testing the results were dramatic, with as much as a 65x improvement in throughput for database replication.
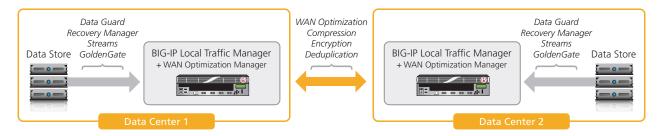

Figure 3: BIG-IP WOM improves throughput on the WAN, speeding replication.

BIG-IP WOM also offloads encryption from the database, which improves not only the performance of replication, but the overall performance of the database itself. Encryption is a CPU-intensive operation that does not have to occur on each server when a BIG-IP device can handle encryption and decryption at the point of necessity. This can help stave off equipment upgrades by freeing CPU processing time for database-centric applications. Moving data into and out of the cloud will play an increasing role in the data center, and encrypting all outgoing data before it enters public space has become all but mandatory for enterprise-class implementations. Offloading that encryption to BIG-IP hardware specifically designed to handle high-volume, large-key encryption will save a lot of processing power on database servers.

While encryption is important, offloading compression to BIG-IP WOM also improves database performance by saving CPU cycles for database processing.

The BIG-IP System and Oracle

The way in which organizations benefit from using BIG-IP products when load balancing Oracle databases varies depending on whether the application infrastructure is a pure Oracle stack (meaning all applications are developed solely using the Oracle client libraries or an Oracle JVM) or a heterogeneous stack (meaning some applications use some non-Oracle development tools).

Pure Oracle Stack		
Feature	Oracle Standalone	Oracle + BIG-IP System
Monitoring	Application SQL Ping	Offload to BIG-IP (one per cluster*)
TCP failover	UCM Connection Pool	Provided by Oracle Net
TCP optimizations	Manual Oracle Net Tuning	Provided by Oracle Net
High availability	Node VIP/Scan IP	BIG-IP system if monitoring enabled
Load balancing	FAN—Runtime Load Balancing	Provided by Oracle Net
Workload management	FAN—Workload Advisory	Provided by Oracle Net
Failure management	FAN messages	Provided by Oracle Net

*One per node in version 10 of TMOS

Figure 4: How the BIG-IP system benefits a pure Oracle stack (a 100 percent
Oracle FAN–capable software architecture).

In the pure Oracle stack scenario, SQL Ping is centralized at the BIG-IP device, with
one or several scripts managing Ping on a schedule best suited to the environment.
Additionally, the BIG-IP system can handle high availability if monitoring is turned on.

Heterogeneous Stack		
Feature	Oracle Standalone	Oracle + BIG-IP System
Monitoring	Application SQL Ping	Offload to BIG-IP (one per cluster*)
TCP failover	Oracle Net Timeout	BIG-IP system, connection reset
TCP optimizations	Manual Oracle Net Tuning	BIG-IP system profiles
High availability	Oracle Node VIP/Scan IP	BIG-IP system, VIP/pool
Load balancing	Oracle Net Connection String	BIG-IP system, instance/name switching
Workload management	Not available	BIG-IP iControl script
Failure management	Not available	BIG-IP iControl script

*One per node in version 10 of TMOS

Figure 5: How the BIG-IP system benefits a heterogeneous stack (not a 100 percent
Oracle FAN–compatible infrastructure).

The benefits of using the BIG-IP system in a load balancing configuration are more
sweeping when there are applications in the data center that utilize database
access methods other than the Oracle SQL libraries. Since "applications" includes
purchased applications, this is the more common scenario. The BIG-IP system
offers all of the functionality that would normally be offered by FAN, and takes
over functions that are not well supported in applications that were not built with
Oracle client libraries.

Conclusion

As workloads continue to increase, organizations will use both load balancing and clustering databases to meet performance goals with commercial, off-the-shelf servers. These methods offer many positive options for database administrators, including high availability through redundancy and load sharing.

F5 BIG-IP products help improve the performance of database clusters by expanding Oracle FAN out to non–FAN enabled clients, thus offering fast connection resets. They also help to load balance non-clustered databases by enabling administrators to bring a database out of production and perform maintenance on it without users noticing that the database is changing. Finally, BIG-IP products help keep remote database replicas synchronized so that shifting load to a replica has a greater probability of success and replication actions take significantly less time, which helps meet RPO and RTO requirements—all while improving performance by offloading encryption and compression.

With databases being such a significant part of the information infrastructure, it is imperative that they be secure, fast, and available. This requires more than just a simple standalone DBMS, and F5 products provide the extra layer to Oracle databases that helps IT management ensure that systems designed around the database are available to users in nearly any circumstance.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119   888-882-4447   www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com