WHITE PAPER: STRATEGIES TO SOLVE AND OPTIMIZE MANAGEMENT OF MULTI-TIERED BUSINESS SERVICES

Symantec.
Confidence in a connected world.

STRATEGIES TO SOLVE AND OPTIMIZE MANAGEMENT OF MULTI-TIERED BUSINESS SERVICES

Niclas Blaback, Technical Product Manager

Table of Contents

Executive Summary

Historically, datacenter administrators have been faced with challenges while managing multi-tier business services. The modern datacenter is often heterogeneous in terms of hardware, software, operating system and virtualization technologies. While new technologies such as virtualization offer cost savings and flexibility, they also introduce complexity in terms of management. Combined with the complexity of a heterogeneous environment, the challenge can be overwhelming. Using different management tools and different management procedures for each tier adds to this challenge.

To address the above issues, Symantec has introduced a new feature called Virtual Business Services. Historically, Symantec has been providing customers with a single product set to standardize storage and availability management in the datacenter. This product set is known as Storage Foundation HA. It provides enterprise class high availability and storage management across multiple platforms, storage arrays and virtualization technologies. With Virtual Business Service and Storage Foundation HA combined together, Symantec provides an end-to-end solution for Storage and Availability management. This white paper covers a set of Virtual Business Services use cases, recommendations and example environments.

Scope of document

The intended use of this document is to provide information about how Virtual Business Services can be implemented. The document describes a number of example problems and how VBS can be used to eliminate them.

Note that this document should not be seen as implementation instructions. This is covered by regular product documentation, which is available on the Symantec Operation Readiness Tool web site (SORT). SORT is located on the following URL: http://sort.symantec.com

Technical support

For technical assistance, visit: **http://www.symantec.com/enterprise/support/assistance_care.jsp**. Select phone or email support. Use the Knowledge Base search feature to access resources such as Tech Notes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Audience

The target audience for this document is datacenter/system administrators, architects, consultants, or anyone who is interested in management and availability of multi-tier business services.

Background

Symantec Virtual Business Services (hereafter known as “VBS”) combines the power of the well-known products VERITAS Storage Foundation, VERITAS Cluster Server, SYMANTEC ApplicationHA and VERITAS Operations Manager.

Together, these products provide an end-to-end solution for High Availability and Management in a multi-tier environment. VBS help customers to efficiently manage applications, high availability and security across multiple tiers, different operating systems and different virtualization technologies.

Introduction

This section introduces the concept of a multi-tier business service and explains how services are deployed in this fashion. In addition, this section covers some common challenges that arise when managing multi-tier business services.

Introduction to multi-tier business services

A few years back, it was common to have all components of a business service running on the same server. For example, if a financial service required a database, an application and a web server, these three components were usually deployed on the same server.

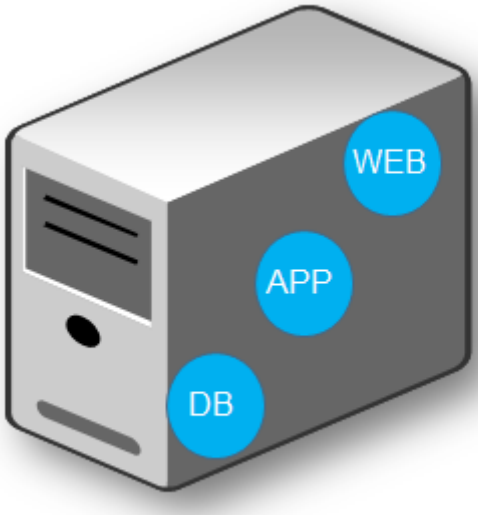

This model has a few advantages, for example:

- Simpler deployment: Everything residing on the same system
- Simpler architecture: Easier to troubleshoot
- Simple cost calculation

As datacenters evolved however, this model also demonstrated some major disadvantages, particularly related to cost and scalability. The disadvantages include:

- Lower hardware utilization: Spare capacity cannot be utilized
- Lack of flexibility
    - Resources cannot be added or removed on demand basis
    - Resources cannot be shared across lines of businesses/departments
- Extensive costs: Same hardware and operating system platform for all components can result in increased costs

Having excessive spare resources in a datacenter is no longer acceptable.
Wasting resources leads to overspending. With shrinking IT budgets and spending cuts, this is simply not an option in today's datacenter.

One way of optimizing the use of hardware resources is to virtualize. Server virtualization is a well-known method to reduce costs in a datacenter. Server virtualization is available on all major platforms today, but it is important to understand that the technologies are hardware specific. As most datacenters are using a heterogeneous mix of operating systems and hardware platforms, it is common to have more than one virtualization technology deployed. Each server virtualization technology comes with different tools and different management procedures.

For example, start and stop of virtual machines are managed differently. In addition, virtual machine operations are usually managed by a virtualization/platform team, and the application is usually managed by an application team. Another pain point can be management of High Availability. If each and every virtualization technology copes with High Availability in a different way, this will also introduce complexity and management overhead. In general, different tools and different management procedures introduce complexity, especially as each technology has to be managed separately.

Virtualized only, physical only, or a combination - regardless of how the underlying hardware is deployed, multi-tier architectures is common in many data centers today. Instead of deploying all components of a business service on the same server, these components are deployed in different tiers. The database is deployed in one tier, the application in another tier and lastly, the web server in yet another tier. This provides some significant advantages, for example:

- Better hardware utilization – services can share hardware
- Cheaper platforms such as Linux and Windows can be used for some tiers, while the I/O intensive components, such as the database is running on larger UNIX servers

- A multi-tier architecture can be fully or partially virtualized or physical only.

In a multi-tier business service, different tiers usually have different requirements. One tier may require full-fledged High Availability with split-second error detection and fast failover, while other tiers just need basic start and stop capability. Hence, the "one-size fits all" approach is not applicable in most cases. The three-tier example business service below, demonstrates the varied requirements and characteristics of each tier:

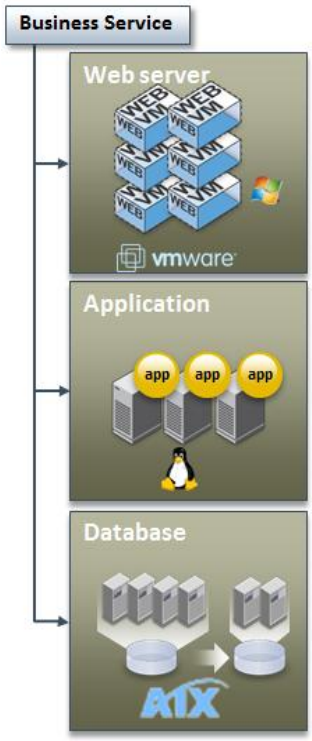

Description of each tier, from the top to the bottom:

- A number of VMWare virtual machines constitute the Web Tier. It is important to ensure that the Web Server inside the virtual machine is online.
- The Application Tier consists of a cluster running on physical nodes. The operating system deployed is RedHat Linux. This application requires application failover capabilities
- The Database Tier is running on AIX. 4 cluster nodes are present on the Primary site, and 2 nodes are present on the DR site. Replication of database data is configured between the sites. This tier requires maximum High Availability, fast failover and tier-1 storage.

The above example business service is a typical multi-tier configuration in many data centers. However, the missing piece is that there is no integration between the tiers. Functionality such as coordinated start/stop, coordinated notification, coordinated recovery, and propagation of failures to remediate issues are missing. The next chapter outlines those tasks and how datacenter administrators have traditionally handled these challenges.

Management of multi-tier business services

Given the advantages of multi-tier architectures, these environments are common in today's data centers. However, multi-tier architectures also introduce a few management challenges. This section describes the challenges and how datacenter administrators manage those challenges.

Business Service Start/Stop

Starting or stopping of a multi-tier business service may seem like a non-issue. In fact, such operations can be complex and time consuming. To understand the challenges, it is important to understand the characteristics associated with the management of multi-tier business services:

- Most multi-tier business services have strict start and stop orders
- The start and stop order is usually managed manually

- The start and stop of each tier needs to be validated. Especially, the start validation can be time consuming
- Different tiers are usually managed by different teams (for example, Windows and Unix). Effective cooperation between those teams is required
- Start/stop procedures can differ significantly across tiers. Different platforms, different passwords, virtualization technologies, hostnames, application stop procedures etc. can make the operation very complex and time consuming
- The actual start/stop command for each tier may be simple, but given the amount of coordination, communication, validation and handover between the different teams, the process can be time consuming
- Start and stop of virtual machines might be required, which may involve yet another team
- Manual procedures can increase the risk of operational mishaps

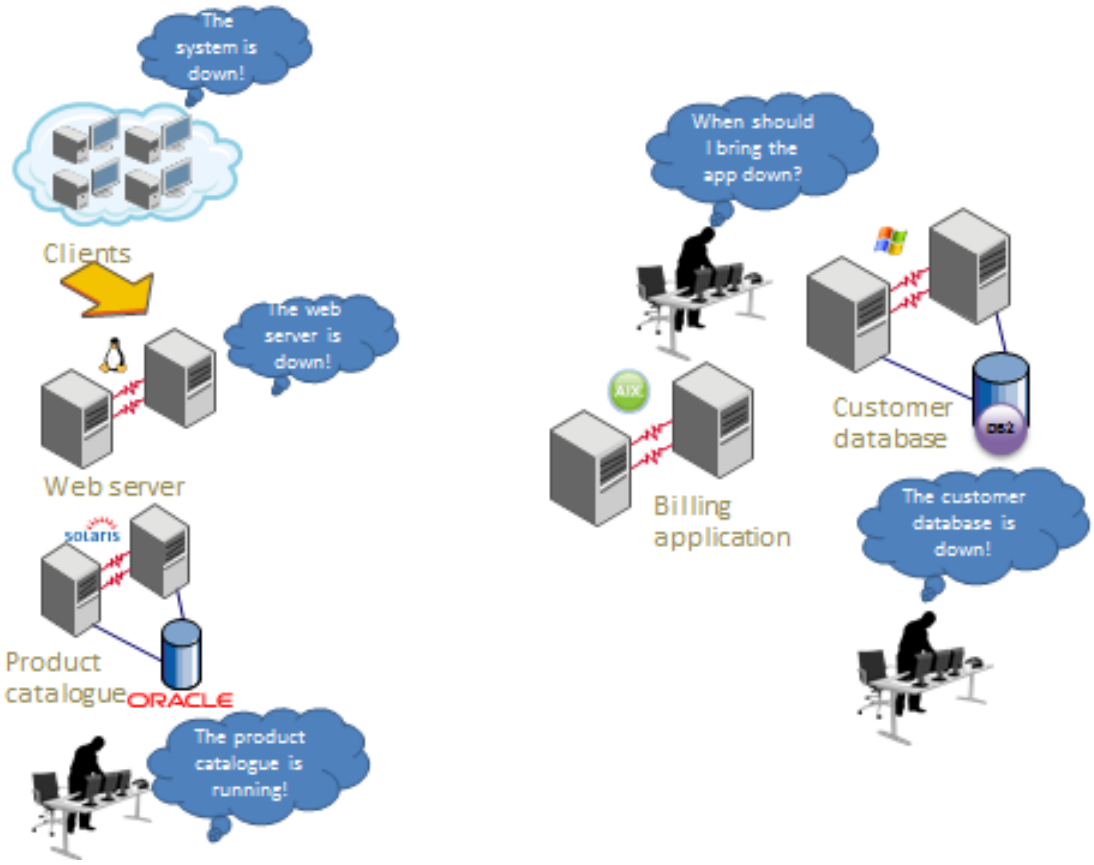

These challenges will be present each and every time the multi-tier business service is started or stopped. This is a regular administrative task in every datacenter. For example, a multi-tier business service may have to be stopped to prepare for a maintenance window. In some environments, datacenter administrators spend a significant amount of time just to login to systems and stop application and/or components. As most businesses operate in a 24/7/365 non-stop environment today, every precaution is made to minimize the time required for a maintenance window. In addition, acceptable planned downtime requirements have significantly gone down.

Even more importantly, these challenges are also present in a real-world Disaster Recovery scenario. When a multi-tier business service needs to be moved from one site to another, all these operations are required to be brought back up at the disaster recovery site as well. Disaster recovery challenges will be covered in the "Disaster Recovery of Multi-tier business service" chapter, later in this document.

Multi-tenancy

One of the key drivers for multi-tier business services is the sharing of resources. Given the shared architecture, multi-tenancy support becomes crucial. For example, a database cluster may be shared among two or more multi-tier business services. In this case, it is very important to make sure that the right set of datacenter administrators have visibility to and can manage the right set of databases, and only those databases.

In the below example, two different Line-of-Business are sharing the same database.

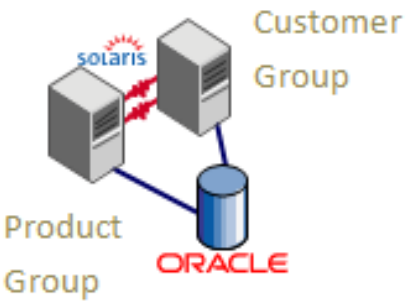

Effective management of a multi-tier business service requires multi-tenancy support. There are a few requirements that have to be met:

- Start/stop of the database has to be managed. For example, the Billing department cannot stop the database without coordinating with the HR department
- Credentials, authentication and permissions have to be managed
- Both departments have to be able to manage their resources, but should not have access to the other department's resources
- Shared resources (in this case the database) should be manageable by both departments

High Availability of multi-tier business services

The sole purpose of a datacenter is to deliver cost-effective and resilient services to business units. From the end user perspective, the availability of the database is irrelevant, if no web server is available to service the request. It is a common mistake to only focus on the availability of the most critical components of a business service, such as the database. While the availability of the database is super-critical, loss of another component of the multi-tier business service renders the database useless. At the end of the day, all components of a multi-tier business service are required to serve the user's request.

Availability of a complete business service

In the introduction section of this document, we outlined that clustering is managed within each individual tier. This will keep the components in each individual tier highly available. However, to ensure availability of the complete multi-tier business service, coordination between the tiers is required. This section describes two example scenarios.

Example scenario 1: Persistent connection between Application and Database

In the two-tier example environment, there is a database cluster and an application cluster. The application running in the application cluster requires a persistent connection to the database. If the database is temporarily offline, the application needs to be restarted to function correctly. Not all applications behave in this manner.
The database is accessed by a virtual IP, floating between the cluster nodes. Initially, the database is running on node "db01" and node "db01" fails due to a hardware issue. The database is failed over to the standby server "db02." The database is now up and running on node "db02."

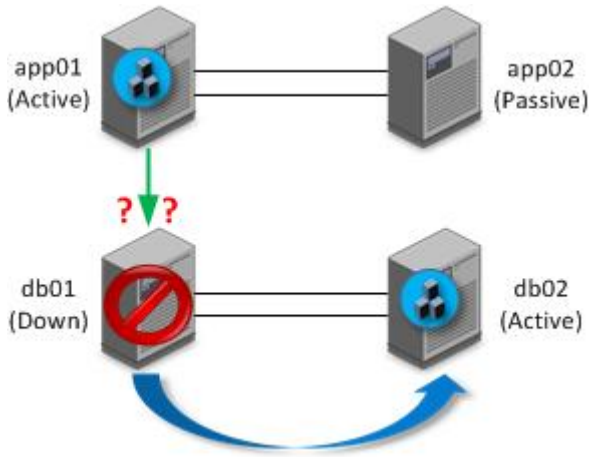

In this case, the database cluster took care of the availability of the database instance. However, as the particular application running in the application cluster requires a persistent connection to the database to function correctly, the application will simply stop working. The application process may be running, but it can't access the data in the database, as the connection to the database has been reset. To remediate the issue, a restart of the application is required.

Some applications are able to reconnect to other components, such as a database, while other applications cannot. This is application dependent, and is not related to other factors such as OS, hardware, clustering software and so on.

Example scenario 2: Application dependent on NFS share

In the two-tier example environment, there is a database cluster and two standalone servers acting as parallel application servers. The first cluster is sharing Network File System (NFS) volumes. The application running in the application servers is dependent on NFS shares from the cluster.

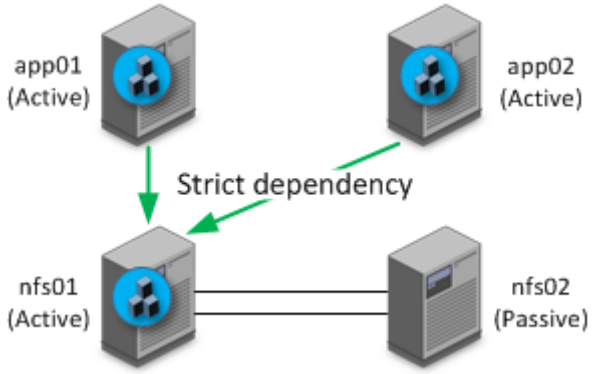

There is a strict dependency between the application and the NFS share. The application cannot start if the NFS share is unavailable. In addition, the application cannot remain online if the NFS share becomes unavailable. In those cases, the application needs to be stopped until the NFS share is back online. When the NFS share is back online, the application can be started again.

Traditional remediation

The two example scenarios above describe two common issues in multi-tier business services. To resolve the example issues, as well as other similar issues, the following remediation methods have been used by datacenter administrators historically:

Manual remediation

Manual remediation is reactive. Commonly, the issue is captured by some kind of monitoring solution (for example based on Simple Network Management Protocol - SNMP). An alert is then sent to the Operations Center. The Operations Center initiates contact with an on-call system administrator. He/she will have to manually login to the system, troubleshoot, find and then remediate the issue.

This method of remediation may work for smaller environments, but for complex environments, it becomes problematic. Manual procedures are time consuming, resource intensive, and also increase the risk of administrative mishaps. In addition, all these manual procedures will increase the amount of downtime, especially as the risk of delay is significant when a large number of human interventions are required. It is important to notice that a manual remediation requires system administrators to be on standby all hours, resulting in an inefficient use of personnel resources and adding to the total cost of ownership.

Scripting
Remediation can also be done by using homemade scripts. The disadvantage of this approach is that it is very costly to manage, require application specific scripting for each tier, password-less access between systems and so on. The other disadvantage is that each scripted solution is very specific to each multi-tier business service. These scripts are also expensive to manage, mainly due to two reasons:

- The scripts need to be sophisticated and intelligent. Basic shell scripts is not sufficient – intelligent decision making and error handling are required
- When the application is upgraded, most likely the scripts needs to be updated as well

Disaster Recovery of multi-tier business services

There is no doubt that a real-world Disaster Recovery scenario is the ultimate test for any business. Disaster Recovery of a complete multi-tier business service can be an operational challenge.
When defining requirements for a Disaster Recovery solution, it is important to understand some of the real-world challenges that a disaster can pose, including:

- Unavailability of key operational team members (such as platform, storage, network etc) on the Disaster Recovery site.
- A complete business service needs to be able to operate from the DR site. It may run on reduced capacity, but it needs to be fully operational. Data replication only is NOT a Disaster Recovery solution
- When considering RPO/RTO requirements, the complete multi-tier business service should be taken into consideration
- Start-up order is important for most multi-tier business services
- Manual procedures, such as changing DNS records, switching replication direction manually and bringing up services in the correct order, is a perfect recipe to break a RTO value. In addition, the whole process needs to be documented in detail
- The failover sequence needs to be coordinated among different teams. Often in sequence due to the start and stop order requirement
- Startup needs to be confirmed by different teams
- Successful regular testing of the Disaster Recovery plan

Declaring a disaster is an exceptional but important decision, usually made under tough conditions. When disaster strikes and the Disaster Recovery plan is initiated, the decision needs to be propagated to all involved teams. As indicated earlier, it is likely that some of the teams won't be reachable in a real-world disaster scenario, which can jeopardize the Disaster Recovery plan.

Organizations usually recognize these challenges during Disaster Recovery testing. It must be noted that Disaster Recovery testing takes place during perfect circumstances, when the event has been long-planned and everyone is well prepared.

In a real-world disaster event, things become even trickier. There are numerous questions that must be addressed such as:

- Are all teams available during the failover process?
- Does coordination between the teams work seamlessly?
- Are key team members able to access the datacenter on-site or remotely? Will that always be the case, regardless of the nature of the disaster?

If the answer to any of the above is "no," then automation is crucial to accomplish a successful disaster recovery.

Limitations of traditional solutions

Datacenter automation solutions

Automating datacenter tasks such as installations, configurations, modifications and sometimes even application start/stop is a common way of optimizing datacenter management. There are several tools available for datacenter automation in the market today.

As the name implies, these tools are used for Datacenter automation and are simply not designed to provide HA/DR capabilities. A key factor to ensure maximum availability is to eliminate any single points of failure (SPOF). These tools are driven from a central orchestration point. One central service that orchestrates tasks in the datacenter is sufficient for day-to-day operations. However, this architecture is not favorable when addressing High Availability and Disaster Recovery requirements. The central orchestration point is a significant SPOF. If the central orchestration point is unavailable (due to any reason such as a corruption, infrastructure issue or a network failure), there is no possibility to initiate the business continuity plan, and get services up and running again. This architectural limitation prevents these tools from being suitable for High Availability and Disaster Recovery purposes.

Symantec Virtual Business Services

To help customers address the issues outlined in earlier chapters of this white paper, Symantec introduced a new feature called Virtual Business Services (VBS).

Symantec VBS combines the power of VERITAS Cluster Server (VCS), SYMANTEC ApplicationHA and VERITAS Operations Manager to provide complete multi-tier business service management and High Availability. VERITAS Operations Manager (VOM) supports the all editions of Storage Foundation, such as Storage Foundation HA, Storage Foundation Cluster File System and Storage Foundation for Oracle RAC.
As VBS now enables management of multi-tier business services on top of VOM and VCS, VOM can now be used as a single tool for Storage and Availability management.

VBS key features include:

- Complete multi-tier management such as coordinated start and stop across different operating systems and/or platforms
- Fault management and propagation between tiers
- Multi-tier Disaster Recovery support, enabling automated Disaster Recovery of a complete Virtual Business Service
- Virtual Machine management support (start and stop)
- Multi-Tenancy and Role-Based Access Control

Virtual Business Services architecture

In the introduction chapter, we outlined a three-tier example business service. Let’s take a look at this multi-tier business service again, and see how VBS can be integrated:

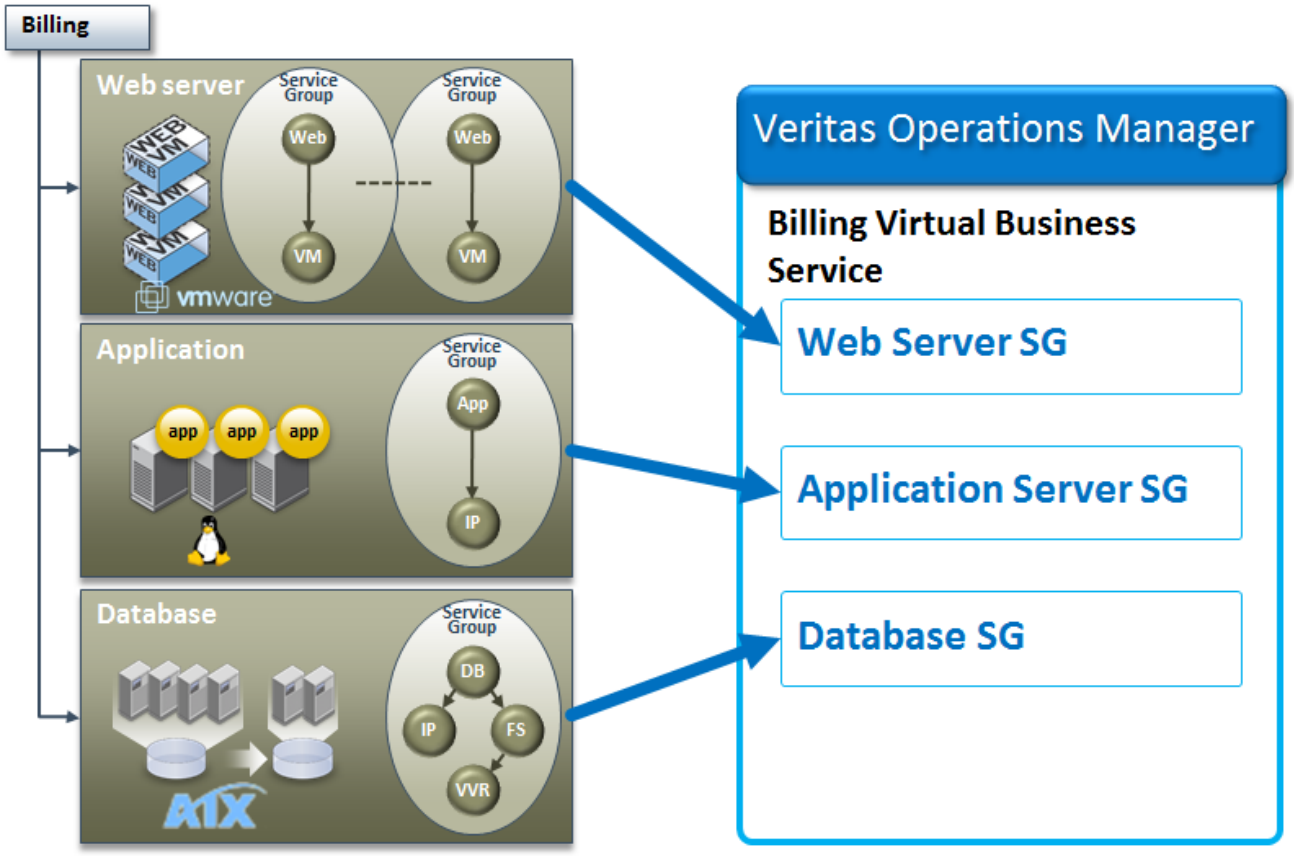

In our example environment, we have the following products installed:

Web Tier:	ApplicationHA 5.1 SP2 for VMWare
Application Tier:	VERITAS Cluster Server 6.0
Database Tier:	Storage Foundation 6.0 HA/DR
VOM CS:	VERITAS Operations Manager 4.1

It is important to understand that High Availability primarily is managed within each tier. The cluster is responsible to keep services highly available within the cluster. The boundaries for an application are the VCS cluster/ApplicationHA instance. Logically, a VBS can be seen as a container that allows service groups to be grouped into a single object. To enable VBS, it is required that each tier has one of the following products installed:

- VERITAS Cluster Server 5.1 or later (including Storage Foundation HA bundles) and/or
- SYMANTEC ApplicationHA 5.1 SP2 or later

Note that VCS and/or ApplicationHA are required in each tier. It is possible to mix and match those products to fit into your environment. In addition, it is required to have at least one VERITAS Operations Manager Central Server:

- VERITAS Operations Manager 4.1 or later

VERITAS Operations Manager is outlined in the next section.

VERITAS Operations Manager and VBS

Veritas™ Operations Manager from Symantec is a comprehensive management platform for Veritas Storage Foundation™ and Veritas Cluster Server environments that optimizes your data center assets. It is a solution to centralize visibility and control, tto ensure availability, to scale operations, to increase storage utilization, and to maintain compliance.

To deploy VBS, it is required to have at least one VOM Central Server in the datacenter. The VOM Central Server is used for configuration, visualization and management of VBS. However, after the initial configuration of a VBS, the VBS can be managed using a Command Line Interface (CLI) as well. VBS does not rely on the VOM

Central Server to operate. CLI operations work regardless if the VOM Central Server is available or not, and the member nodes of a VBS will operate autonomously of the VOM Central Server once VBS is deployed.

More information about VOM, download links and white papers can be found on the following URL:
http://go.symantec.com/vom

Managing multi-tier business services using VBS

A key architectural feature of VBS is the ability to define dependencies between Service Groups. Primarily, VBS dependencies are used for two reasons: Start/stop ordering, and fault propagation.

Historically, VCS has been providing dependencies between service groups. However, there is a key difference between VCS Service Group dependencies and VBS Service Group dependencies:

- VCS Service Group dependencies are configured between service groups **within** a single cluster, and therefore on the same platform
- VBS Service Group dependencies are configured between service groups in **different** VCS clusters and/or ApplicationHA instances. VBS Service Group dependencies are supported across different platforms and operating systems

Note that it is possible to combine VCS and VBS Service Group dependencies. VBS applies a specific logic to Service Groups that has VCS dependencies configured.

Following the same model as VCS, VBS is also using the parent/child model for service group dependencies. The following picture describes the parent/child model:

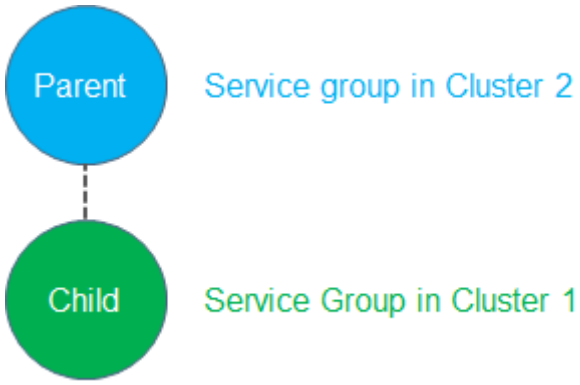

Note the difference between VBS dependencies and VCS dependencies. The above picture displays a VBS dependency as the service groups resides in different clusters.

Start/Stop Ordering (soft dependency type)
Virtual Business Services enables coordinated start and stop of service groups in multiple tiers. Instead of coordinating start/stop activities among different platform teams/application teams, a datacenter administrator can start or stop a complete VBS in a single operation. The soft dependency is used to define start and stop order.

Fault propagation between VBS dependent service groups is not enabled when using the soft dependency VBS fault propagation is outlined in the next chapter - High Availability of multi-tier business services using VBS**.**

VBS provides the ability to start and/or stop VMWare Virtual Machines during VBS start/stop operations. In a single operation, a complete VBS can be started or stopped, including any virtual machines. If a Virtual Machine is in a stopped state, it can automatically be started by VBS. VBS tracks the start-up and also ensures that the application running inside the Virtual Machine is started correctly. The same logic applies for the stop of the VBS.

This is especially useful in Disaster Recovery use cases. Virtual Machines on the Disaster Recovery site can remain in stopped state, and started when they are required. This will happen in an automated, but controlled fashion, as the dependencies are predefined.

VMWare Virtual Machine start/stop is supported with the following VBS combination:

- Supported version of ApplicationHA (5.1 SP2 or later) running inside the VM

- VOM Control Host Add-on 4.1 or later (can be downloaded free of charge on http://go.symantec.com/vom)

The application inside the Virtual Machine is required to be under ApplicationHA control. Further releases of VBS may support VM start/stop for other virtualization technologies.

Visibility and overview
VERITAS Operations Manager allows you to centrally manage application, server, and storage environments from a secure Web-based console. You can also perform wizard-driven operations on file systems, volumes, and disk groups with a single right-click. With a summary dashboard that displays the state of these resources, administrators are able to monitor the data center infrastructure and drill down to manage any risks or faults in the environment.

Using the VOM Web-based console, VBS offers full visibility and displays the consolidated health of a complete multi-tier business service. Instead of examining each individual tier, VOM provides the overview, and can also send alerts in case an issue is encountered.

In addition to the VOM web-based console, VBS also provides a powerful CLI which enables administration and visibility into the state of a complete multi-tier business service.

Multi-Tenancy using VBS

VBS supports holistic grouping and management of resources.
This is handled using a VOM feature known as Business Entities. VBS is built upon Business Entities and extends the functionality to provide HA/DR capabilities.

The VOM access model supports RBAC (Role Based Access Control) that allows different Lines of Businesses to control their resources. This model provides isolation between VBSs, and it also supports shared resources.
In the environment below, three VBSs are configured; Billing, Finance and HR.

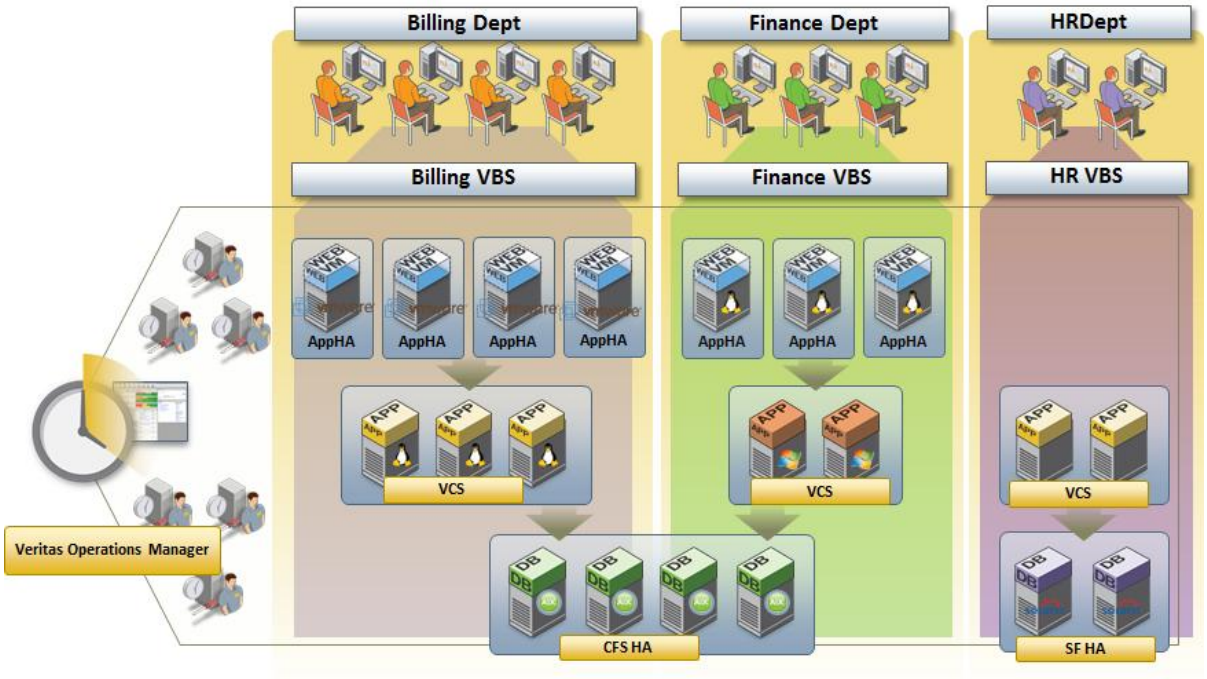

The Billing department has full control over the Billing VBS. The same goes for the Finance VBS and the Finance department. However, the database tier (bottom-most) is shared between these two departments. The VOM access model makes sure that the resources can be separated. Depending on the environment, multi-tenancy can be configured in two different ways:

- Each Line of Business has a separate database. In this case, the Billing database and Finance database will be placed in two different VCS service groups in the database tier. There will be no service groups shared between the two VBSs.

Or:

- The Billing and Finance Lines of businesses are utilizing the same database VBS introduces a concept known as shared VBSs. In this case, the database is shared. A service group can be shared among two or more VBSs.

Using the two models, VBS provides flexibility for shared infrastructures. Objects can be shared, as well as isolated between different VBSs.

High-Availability of multi-tier business services using VBS

When VCS and ApplicationHA manage High Availability in each tier, VBS takes care of the availability of the complete multi-tier business service. VBS allows datacenter administrators to configure dependencies with fault propagation policies between service groups running in different tiers, providing an even higher level of availability for the entire multi-tier business service.

Inter-cluster fault propagation (firm and restart dependency types)

The firm and restart dependency types includes the same functionality as the soft dependency type (start/stop ordering, VMWare VM start/stop), but it also includes a feature known as "inter-cluster fault propagation."

It is important to understand that inter-cluster fault propagation is triggered on faults only. For example, if a failover occurs in a lower tier, an event can be propagated to another tier to automatically remediate an issue. If an administrator is issuing a switchover, that is considered as a controlled, manual operation, and VBS does NOT propagate events to an above tier.

NOTE: When VMWare virtual machine start/stop is configured, it is important to understand that the virtual machine itself is not affected by the fault propagation process. VBS will only start/stop a virtual machine during a VBS start/stop operation. Fault propagation will affect the service group running inside the virtual machine, but not the virtual machine itself.

Restart Dependency

When a child service group faults, the parent service group ignores the event. When the child service group recovers, the parent service group is restarted.

In the "Availability of a complete business service" section, Example Scenario 1, we outlined an example where an application needs to be restarted to remediate an issue by using the VBS restart dependency, the remediation can be triggered automatically.

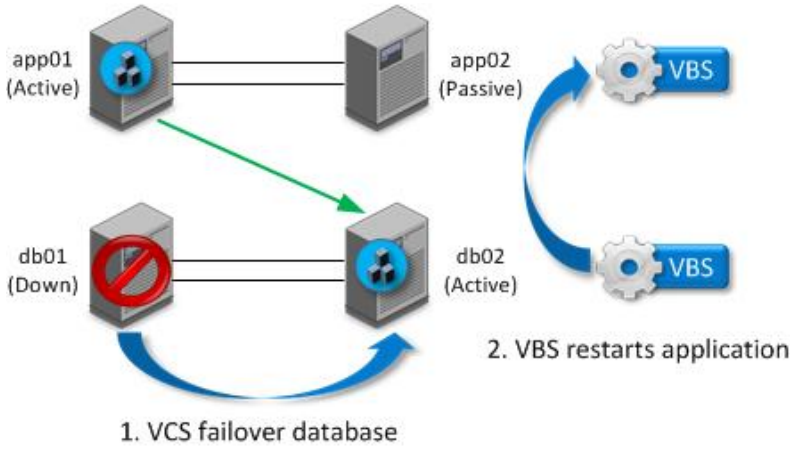

In the above configuration, the following process occurs:

1. VCS detects that node db01 is down and performs a failover of the database (and associated VIP) to node db02.
2. VBS detects that the database failed over to db02. As the database service group is the child service group, VBS propagates the failure and triggers a restart of the application service group.
3. When the application service group is restarted, a new connection will be opened to the database, allowing the multi-tier business service to operate again.

Firm Dependency

When a child service group faults, the parent service group is taken offline. When the child recovers, the parent is brought online.

In the "Availability of a complete business service" section, Example Scenario 2, we outlined an example where the application needs to be taken offline when the NFS share is offline, and then online when the NFS share is online. By using the VBS firm dependency, this process can be automated:

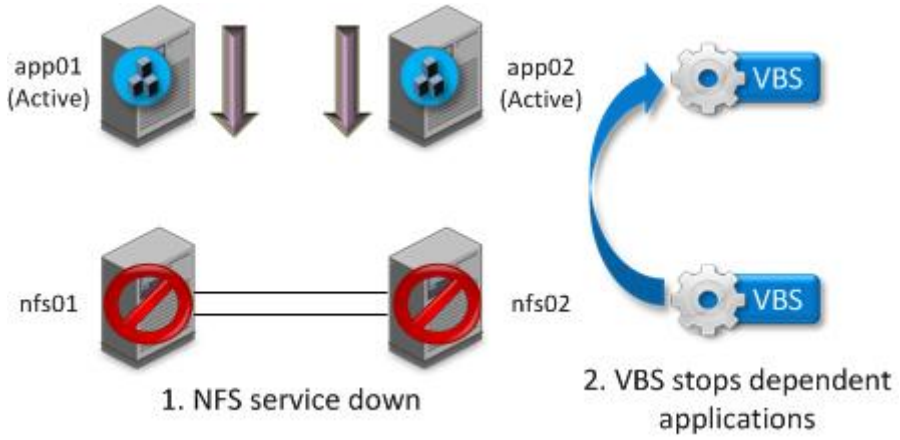

1. The storage connectivity on all nodes in the bottom-most cluster fails due to a zoning issue. Root cause for the event is an administrative mishap.
2. As a VBS firm dependency is configured between the NFS sharing service group and the application service groups, all application service groups are brought offline

When the storage issue has been resolved, and the bottom-most cluster is able to online the NFS service again, the following will happen:

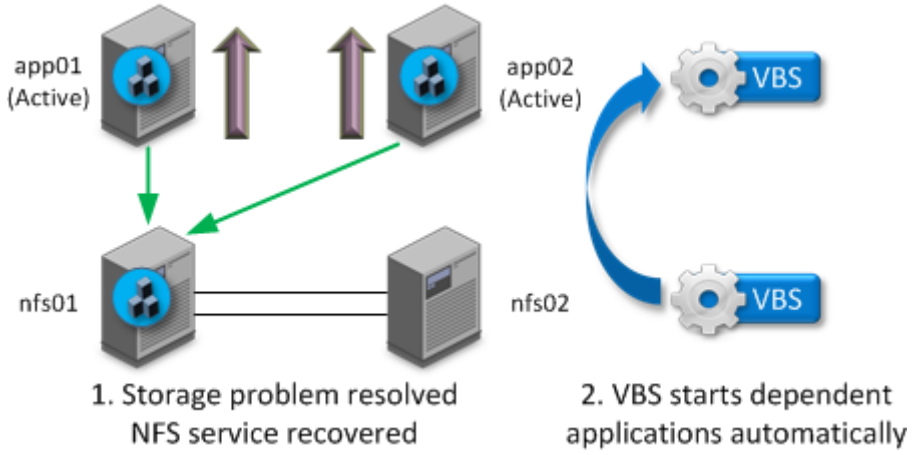

1. VCS is starting the NFS sharing service group. VBS remembers the previous failure, and starts the application service groups. The multi-tier business service is operational again.

Dangling service groups

Service groups also can be added into a VBS without any dependencies. This is known as dangling service groups and is used for start/stop coordination only. It is especially useful in Disaster Recovery configurations. In addition, one VBS may have more than one dependency trees, or a combination of dependency tree(s) and dangling service groups, or only dangling service groups. Below is an example of a VBS containing dangling service groups:

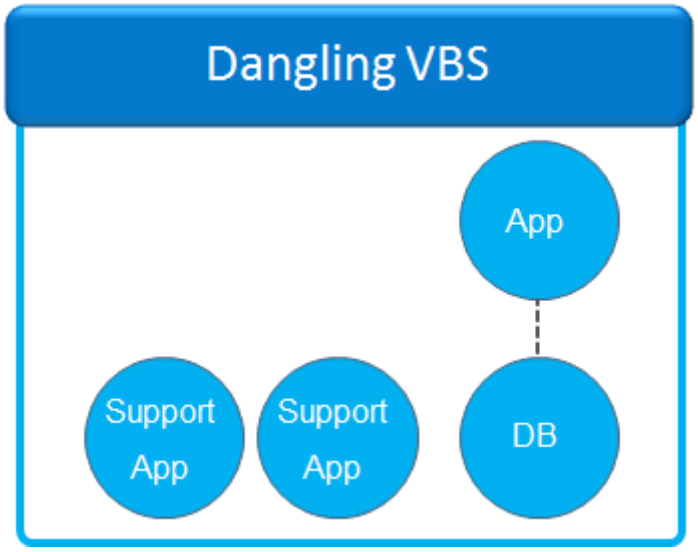

Disaster Recovery of multi-tier business services using VBS

By automating the whole process of bringing up a complete multi-tier business service, VBS ensures that the start order is honored. VBS validates that each and every tier starts up correctly. Hence, VBS together with VCS and/or ApplicationHA can ensure complete multi-tier business service startup at the DR site.

The underlying components, such as storage replication, application start procedures, changing DNS records and so on, are taken care of by VCS and/or ApplicationHA. As the VBS start/stop procedures are predefined by using the dependencies, it ensures that the multi-tier business service is started in the correct order, and with the correct procedures. By using VBS for this purpose, manual procedures can be eliminated.

The administrator will simply have to initiate failover of the VBS object from VOM or CLI to bring up the entire VBS on the Disaster Recovery site. All required procedures are predefined, and VBS takes care of those procedures the background.

When deploying VBS for Disaster Recovery, there are some design considerations including:

- Each individual multi-tier business service requires one VBS per datacenter. Two datacenters are available on our example environment, hence we will have two VBS's
- Disaster Recovery failovers with VBS is fully automated and always triggered manually
- It's recommended to configure one VOM Central Server per site.

The picture below displays a sample VBS Disaster Recovery configuration.

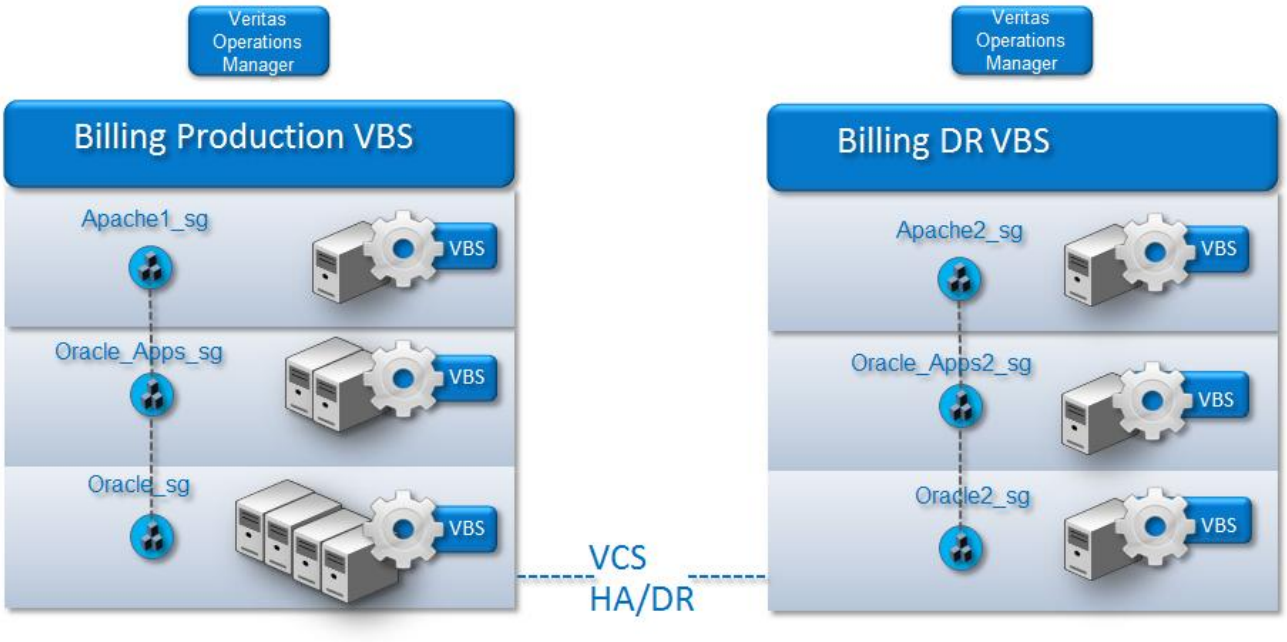

When VBS is deployed in a Disaster Recovery fashion, two VBS objects are required - one VBS per site. In this example, Billing Production VBS and Billing DR VBS are configured. Note that number of nodes etc. differs between the Production site and the DR site. VBS supports different configurations on the Production and Disaster Recovery sites. The number of tiers can differ as well.

Requirements for VBS Disaster Recovery:

- At least one tier needs to be installed with VCS HA/DR, which includes Global Cluster Option (GCO). GCO needs to be configured between the sites

Note: In our example, GCO is configured between the bottom-most tier. This is the most common configuration, however the requirement is that GCO is configured for least one tier – it can be any tier participating in the VBS.

Example deployments

This section provides a number of example VBS deployments. The purpose is to assist customers with VBS design and architecture, in preparation for an implementation. It is important to understand that data centers are different. The below environments are just examples. Many of these examples can be mixed and matched to fit into your datacenter environment. The last section of this chapter contains a comparison between all the Example deployments.

Example I - Two-Tier Business Service (database and application)

This example environment consists of a four node VCS database cluster running on Solaris and a two node VCS application cluster running on Linux Redhat. The business service is used by a financial department and the VBS name is finance VBS.

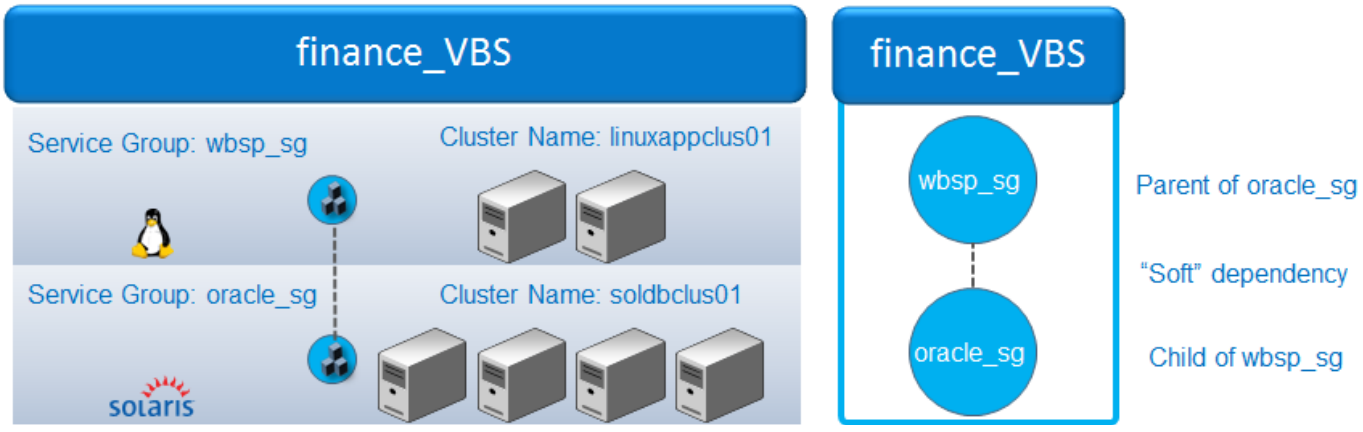

Example II - Three-Tier Business Service (database, app, web). App and web Tiers virtualized.

This example environment consists of a four node SFHA database cluster running on HP-UX, a two node VCS cluster running on Redhat Linux. In the top, a webserver tier consisting of eight VMWare virtual machines, running Windows and IIS under ApplicationHA control is present. In this example, fault propagation is configured between the database cluster and the application cluster. The reason for this requirement is because the application requires a persistent connection to the database.

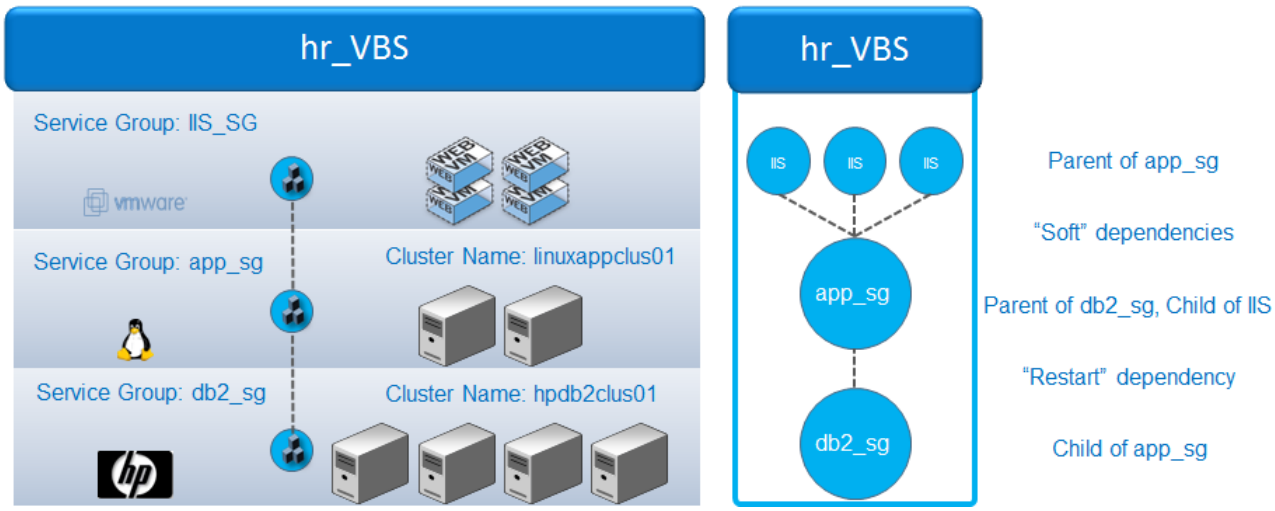

Example III - Two-Tier environment with a shared database (database, two apps)

In this environment, two different applications are deployed in the application tier. These two applications are accessing the same database. The database service group is a shared service group. The database cluster is running on AIX, and the application cluster is running on Linux.

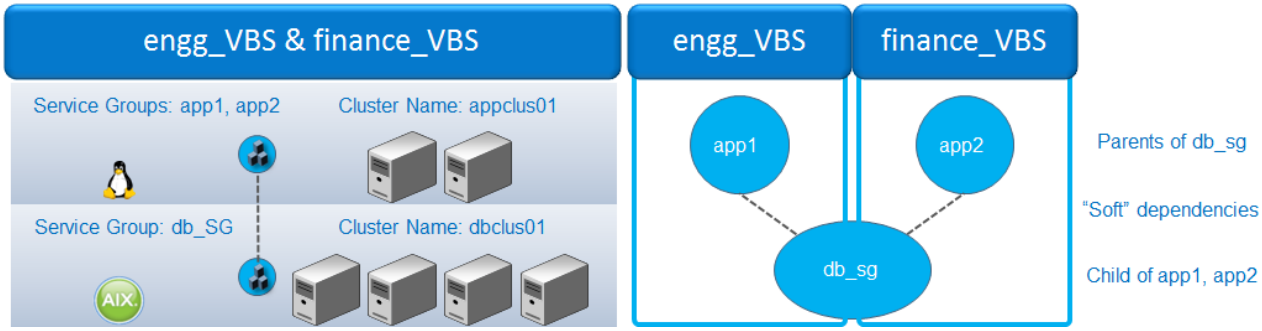

Example IV - Three-Tier Business Service with DR (database, app, web)

This example environment consists of three different tiers. The first tier is a database tier running on Solaris. This tier has four cluster nodes on the production site and two cluster nodes on the Disaster Recovery site. These two clusters are connected by using VCS HA/DR (Global Cluster Option).

The second tier is an application tier configured for local High Availability. 2 by 2 node clusters are deployed, which means one application cluster per site.

The third tier is a web tier running in virtual machines. On the production site, there are four virtual machines available. The DR site hosts two virtual machines. On the DR site, VBS is configured for VMWare VM start/stop, which means that the VMs will be powered off when the VBS is stopped, allowing resources to be used for other purposes.

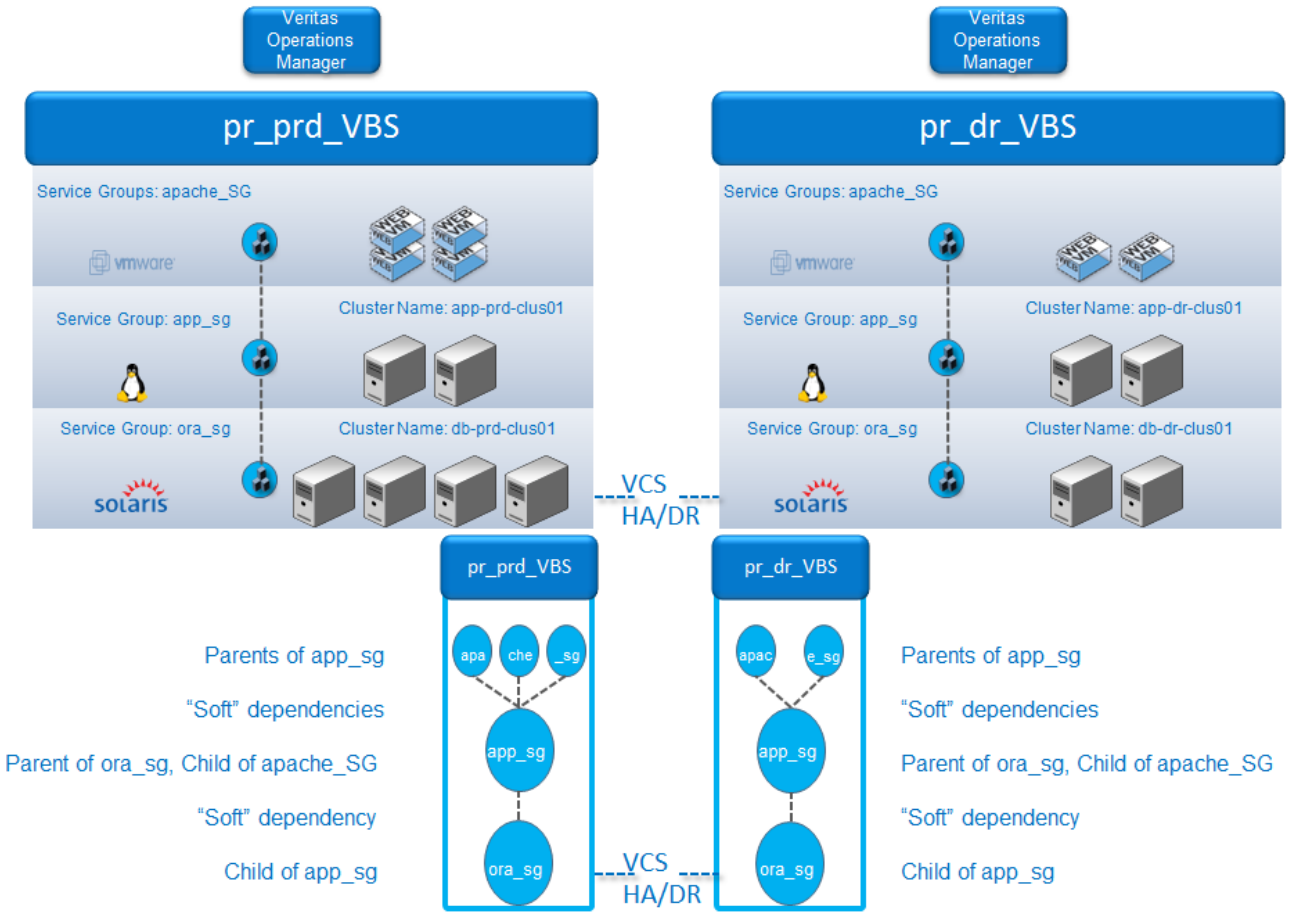

VBS Functionality comparison among the Example Deployments

The table below displays a comparison between all example deployments and the VBS functionality that is used in each case.

VBS feature	Example I	Example II	Example III	Example IV
Multi-tier Visibility	X	X	X	X
Multi-tier Management	X	X	X	X
Multi-tier HA (fault management)		X		
Multi-tenancy by using Role Based Access Control	X	X	X	X
Shared Service Groups			X	
VMWare VM start/stop				X
Multi-tier Disaster Recovery support				X

Below is an explanation of each feature:

- Multi-tier Visibility – this feature provides full visibility and status reporting from the VERITAS Operations Manager console. This simplifies troubleshooting and provides a detailed overview about the health of the multi-tier business service
- Automated, single-click start/stop procedures for the entire multi-tier business service
- VBS Fault propagation between clusters remediates application-specific issues automatically, eliminating manual procedures, thus providing an even higher level of availability
- Role Based Access Control security model, supporting multi-tenancy
- Multi-tenancy support by using shared service groups
- VMWare VM start/stop as a part of VBS start/stop operations
- Fully automated Disaster Recovery, which eliminate manual steps in the Disaster Recovery plan

Conclusion

Today's datacenters have to constantly evolve to meet the challenges of a business environment that needs to operate in a 24/7 environment. As a result, data center administrators have to manage multi-tier business services to accommodate the business needs. Virtual Business Services resolves a number of issues that datacenter administrators are constantly faced with. These issues include providing high availability and managing such a complex environment. Virtual Business Services resolves these challenges through automation and complete visibility and it is easily integrated on top of existing infrastructure.

About Symantec Symantec is a global leader in providing security, storage and systems management solutions to help businesses and consumers secure and manage their information. Headquartered in Mountain View, CA., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site.
For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
350 Ellis Street
Mountain View, CA 94043 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

Last updated April 2012