



connect
OpenVMS
Bootcamp 2010

HP OpenVMS ServiceControl

Making non-cluster aware applications highly available

Dr. Wolfgang Burger, HPS Technology Services/Austria

Agenda


- Introduction
- Requirements
- Basic Concepts and Terminology
- Controlling OSC behavior
- Break
- Hands-on session
 - Configuration from scratch
 - OSC behavior on various failure conditions

Introduction

- HP OpenVMS ServiceControl (OSC) is a management framework that makes non-cluster aware services (applications) highly available to their clients.
- Basic Functionality
 - Failover
 - Switchover


Basic Functionality

- Failover



Basic Functionality

- Switchover



Requirements

- Monitor the health of a service
 - Application plus all its required hardware and software resources (disks, shadow sets, network interfaces ...)
- Relocate services due to service or system failures
- Service switchover on user request
- Able to handle any kind of application
 - Easily extendable
 - Reuseability of exiting Management scripts (Startup, Shutdown, Monitor Scripts)
- Easy to configure and control
- Event notification
 - Ordered by occurrence
 - Freely definable event notification method

Requirements (contd.)


- Common management interface
- Common configuration interface
- Terminology and management interfaces similar to VERITASTM cluster server (VCS)
- Wide range of control attributes to adjust OSC behavior to customer needs
- GUI to
 - Visualize actual state of the services and resources
 - Visualize resource and service dependencies
 - Manual control and management of the services and resources

Basic Concepts and Terminology

- OSC cluster
- OSC management entities (building blocks)
 - Resource
 - Service
 - Service Group

What is an OSC cluster?

- Consists of all or a subset of OpenVMS cluster members



OSC Cluster

- OSC cluster members have to be defined when OSC is configured
 - Votes have to be assigned to each OSC cluster member
 - Quorum calculation (done the same way as OpenVMS cluster)
 - OSC quorum is different from OpenVMS quorum
 - As long as the OSC cluster has quorum the system resources of all remaining OSC cluster members are sufficient to run all managed applications
 - When OSC quorum is lost due to a node failure:
 - Applications that had been online on the failing node are not automatically started on the remaining OSC cluster member, since an OSC quorum lost state signals that overall system resources are not sufficient to run all managed applications
 - OSC blocks interactive management commands except:
 - » SHOW and ADJUST QUORUM commands


OSC Cluster (contd.)

- Special Quorum handling on a 2 node cluster with a quorum disk
 - Quorum disks cannot be configured as a voting member
 - OSC would block if one node fails although the OpenVMS cluster is still up.
 - OscCtrlAutoAdjustQuorum cluster attribute
 - If this attribute is set to TRUE, OSC automatically adjusts (OSC) quorum whenever an OSC cluster member unexpectedly leaves the OpenVMS cluster.

Resources

- Resources are hardware and software entities such as:
 - Disks
 - Shadow Sets
 - Network Interface Cards
 - IP Addresses
 - Databases
 - Applications
 - ...
- Resource dependencies define the order in which resources are brought online or taken offline
 - Child resources must be online before a parent resource can be started
 - Parent resources have to be offline before a child resource can be shutdown

Resource Dependency Example



Resource Criteria

- An OSC resource is whatever you define as a resource
- Nevertheless, an OSC resource has to fulfill criteria in order to be managed by OSC
 - An OSC resource must be capable of being explicitly started by a set of commands
 - Startup of one OSC resource must not implicitly start other OSC resources. If so, these resources cannot be defined as independent OSC resources.
 - Each instance of a resource type (i.e. database) must be capable of being stopped without affecting other instances of that resource type.

Resource Criteria


- An OSC resource must be capable of being explicitly stopped, by forcible means if necessary.
 - Shutdown of one OSC resource must not implicitly stop other OSC resources. If so, these resources cannot be defined as independent OSC resources.
- OSC resources must be capable of being monitored.
- OSC resources must be crash tolerant.
 - An OSC resource runs on a node that crashes -> OSC resource will start on a failover node in a none state (i.e. no memory content required to start the resource)
- OSC resources must be host independent within an OpenVMS cluster.
- No license restrictions or host name dependencies that prevents successful failover.

Resource Categories

- On-Off
 - Monitored
 - Started if required
 - Stopped if required
 - *Eg.* IP-Address
- On-Only
 - Monitored
 - Started if required
 - *Eg.* Shadow Set
- Persistent
 - Monitored only
 - *Eg.* Network Interface Card

Service

- A service is a logical grouping of resources and resource dependencies that are required to run a dedicated service (application). It is the management unit that controls resource sets .




Service Dependency

- Service dependencies define the order in which services are brought online or taken offline
 - Child Services must be online before a Parent Service can be started.
 - Parent Services have to be offline before a Child Service can be shutdown.

Service Group

- A service group is a logical grouping of services and service dependencies. It is the management and failover entity of OSC.



Service Group


- Service Group Types
 - Failover
 - A failover service group runs on one system in the OSC cluster at a time. Failover groups are used for non OpenVMS cluster aware applications (ie. applications that are not designed to maintain data consistency when multiple copies are started).
 - Multi-Instance
 - A Multi-Instance service group is active concurrently on more than one but not on all systems within the OSC cluster. All services within the service group must be cluster aware.
 - Parallel
 - A parallel service group runs concurrently on all OSC cluster members. All services within the service group must be cluster aware.

Execution node list and priority has to be defined for each service group

OSC components

- OSC Agents
- OSC Service Engine
- OSC Master Engine
- OSC Management Utility
- OSC Configuration Utility
- OSC Configuration Database
- OSC Event Notification Service

OSC software architecture



OSC Agents

- Resources of a particular type are managed by one OSC Agent (i.e. Shadow sets, disks)
- Resource management means
 - Monitoring the status of a resource
an OSC agent decides whether a resource has failed
 - Starting a resource
 - Stopping a resource
 - Cleaning up a resource
- OSC Agent consists of the OSC Agent framework and action routines that provide the resource type specific logic
- OSC Agent framework provides processing logic
 - Understand common resource attributes
 - Workflow logic
 - Communication with the OSC Service Engine

OSC Agent action routines

- Open
 - Initializes the resource
 - Called once when the agent starts monitoring the resource
- Monitor
 - The monitor action routine is periodically called to determine the resource state and to verify whether the resource state has changed. The online and offline monitoring interval can be configured resource specific.
 - The monitor action routine is called for all managed resources whenever the OSC agent is re-started and after every attempt to put a resource online or offline in order to verify that the operation was successful.
- Online
 - The online action routine brings a specific resource online from an offline state.

OSC Agent action routines (contd.)

- Offline
 - The offline action routine shuts down online resources
- Cleanup
 - The cleanup action routine is called (forced shutdown) for a resource after a resource has failed to come online, failed to go offline, or failed while in an online state. The cleanup action routine has to be designed to forcibly shutdown a resource when it has failed in order to ensure that the resource does not remain in an undefined state. The cleanup action routine will be executed only for On-Off resources, since these resources are typically not cluster aware. On-Off resources have to be offline before they can be brought online on another OSC cluster member.

OSC Agent action routines (contd.)

- Actions routines can be implemented either using DCL scripts or as C functions
- OSC Agent action routines can be defined resource specific
 - Thus, an OSC Agent may call different action routines for resources of the same type
- Due to this design it is easy to develop new OSC Agents

OSC Service Engine

- Active on all OSC cluster members
- It monitors and controls the OSC agents on a node:
 - Starts all the required OSC Agents on a node whenever the OSC environment is started on a particular node.
 - Stops all required OSC Agent processes when the OSC environment is shutdown on a node.
 - Guarantees that all pending service group, service and resource transactions complete before the shutdown request is executed.
 - Monitors the status of the OSC Agents running on the node. The OSC Service Engine periodically checks the receipt of heartbeat signals from the OSC Agents. If this check fails (it has received no heartbeat message from an agent within a predefined time interval) for an OSC Agent, the OSC Service Engine automatically restarts the appropriate OSC Agent if it is allowed to. The agent fault processing behavior by the OSC Service Engine can be defined agent specific.

OSC Service Engine (contd.)

- Maintains/updates the status of services and service groups locally configured according to the status information received from the OSC Agents.
- Forwards the resource, service and service group status information to the OSC Master Control Process.
- Guarantees that all resource, service and service group administrative commands received from the OSC Master Control Process are executed according to the configured dependencies.
 - E.g. if the OSC Service Engine receives the online command for a service group from the OSC Master Control Process, the OSC Service Engine ensures that all services of the service group and all resources defined within each of these services are brought online bottom-up according to the configured dependencies.

OSC Master Control Engine

- Started on all OSC cluster members
- Active only on one OSC cluster member
- All standby OSC Master Control Engine instances check the health of the active OSC Master Control Engine
- Knows the status of all service groups on all OSC cluster members. Thus, the active OSC Master Control Engine is the one that decides whether to put a service group online, offline or failover .
- The OSC Master Control Engine checks the health of the OSC Service Engines running on the OSC cluster members
 - Checks the heartbeat signals sent by the OSC Service Engines.
 - If the OSC Master Control Process does not receive a heartbeat within a predefined time period from a particular OSC Service Engine it automatically tries to restart that OSC Service Engine according to the OSC Service Engine control parameters

OSC Master Control Engine

- The active OSC Master Control process provides the console interface for interactive OSC management.
 - Up to 64 console links are supported
- SCS layer utilized for reliable communication between the OSC Service Engines

Failover Policies

- Failover Policy is OSC cluster wide defined
- Static
 - Defines a static node failover list for a service group (application)
- Load-Balancing
 - Distributes the configured service groups among the cluster members depending on the free load capabilities on OSC cluster members
- Service Group exclude lists
 - Can be defined in addition to the failover policies defined above
 - Defines the service groups (applications) that are not allowed to run on the same node.

Special Resource Type

- Cluster Locked Resources
 - A cluster locked resource is a resource that will be started only on one node within the OSC cluster regardless if the service group that owns this resource is configured to run concurrently on different nodes
 - Service Group exclude lists
 - Cluster Locked Resource exclude list
 - Can be defined for any Cluster Locked Resource
 - Defines the Cluster Locked Resources that are not allowed to run to the same node

OSC Management Utility

- OSC\$BIN:OSC\$MGR.EXE
- Automatically connects to the active OSC Master Control Engine
- Management actions:
 - Status display of the OSC cluster environment
 - Status display of all service groups, services and resources
 - Management of all service groups, services and resources
 - Setting a service group online, taking it offline
 - Service group switchover
 - Freeze / disable a service group or a resource
 - Clear faults and change ADMIN_WAIT state
 - ...
 - Management of the OSC cluster environment

OSC Configuration Utility

- OSC\$BIN:OSC\$CFG.EXE
- Common OSC Configuration Utility to
 - Manage OSC configuration projects (databases)
 - Configure within an OSC configuration project
 - OSC cluster
 - OSC Master Engine control parameters
 - OSC Service Engine control parameters
 - OSC agent definition and control parmameters
 - Resources and resource dependencies
 - Services and service dependencies
 - Service groups
 - OSC events and notification method

Basic OSC behavior on resource faults

- Basic actions performed by OSC when a resource fails
 - Managing resource faults enabled?
 - Cleaning up Resources
 - Critical or Non-Critical Resources?
 - Fault Propagation enabled?
 - Service Priority
 - Service Group Failover

Basic OSC behavior on resource faults (contd.)

- When does OSC consider a resource as faulted?
 - OSC considers a resource as faulted when:
 - the resource state changes unexpectedly.
 - Eg. an online resource goes offline.
 - a required state change does not occur.
 - Eg. a resource failed to go online or offline.
 - Typically OSC Agents take predefined actions to correct an issue before reporting a resource fault condition to the OSC Service Engine.
 - E.g. the agent may try to bring a resource online several times before reporting a fault.

Basic OSC behavior on resource faults (contd.)

- Managing resource faults enabled?
 - OSC takes automatic actions on a resource fault only if the service group of the failed resource is configured to manage faults
- Cleaning up Resources
 - When a resource fails and fault management is enabled, OSC calls a cleanup action routine to forcibly shutdown the resource. If the cleanup action routine fails to shut down a failed resource the resource state is undefined and so service group failover processing will not be initiated. Failover of a service group that contains undefined resources may cause concurrency violations.
 - The cleanup routine is only called for Off-On resources since these resources are - in OSC terminology - not cluster aware. All other resource types (On-Only, Persistent) are cluster aware resources that can be online concurrently on multiple nodes. Thus, a service group failover due to a fault of an On-Only or Persistent resource never causes concurrency violations.

Basic OSC behavior on resource faults (contd.)


- Critical and Non-Critical Resources
 - The Critical attribute of a resource defines whether or not a resource fault initiates fault propagation to the service group level.
If a resource is configured as non-critical (by setting the Critical attribute to FALSE), services and service groups that contain this resource will not be marked as faulted. Failover processing will be initiated only if the service group is declared as faulted. Thus, faults on non-critical resources will not cause service group failover.
- Fault Propagation
 - Fault propagation is the name given to the process of marking a service and a service group as faulted when a resource fails.
 - If fault propagation is disabled the services and service groups of the resource will not be marked as faulted and thus, the service groups will not failover.

Basic OSC behavior on resource faults (contd.)


- Service Priority
 - Once a service is marked faulted, OSC checks the service priority of all independent service trees configured within the service group. If the service priority of an online service tree is greater than the value of the failed service tree, no service group failover is initiated.
- Service Group Failover
 - If the service group is marked as faulted it takes the whole service group offline
 - If service group failover is enabled On-Off and Multi-Instance, service groups are started on another node within the OSC cluster providing a valid node to start the service group exists

OSC behavior diagrams


- Symbols used to illustrate OSC functionality when a resource fails




Critical Resource fails




Non-Critical Resource fails




Critical Resource fails to come online



Critical Resource fails – fault propagation is disabled



Critical Resource fails – fault management is disabled



OSC Event Notification

- OSC provides event notification messages when:
 - A resource, service or service group state changes
 - The OSC cluster state changes
 - The operational state of any component of OSC changes
(OSC Agents, OSC Service Engine, OSC Master Control Process)
 - Any of the OSC cluster members change state
(OSC cluster member leaves or joins the OSC cluster)
- OSC executes automatic or user initiated administrative operations to:
 - Recover resource fault conditions
 - Set service groups Online/offline/failover/switchover
 - Reestablish OSC cluster integrity
(restart of operational OSC components, removing a member from the OSC cluster if a node fails ...)

OSC Event Classes

- Nine OSC Event Classes exist:
 - HEARTBEAT
 - Heartbeat events are triggered whenever an OSC component receives a heartbeat message from one of its managed OSC components.
 - OSCAGT_CONTROL_EVT
 - This OSC Event Class contains resource control messages. Resource control messages are triggered by OSC agents when they execute administrative operations (like online, offline commands) either on request by the OSC Service Engine or triggered automatically to recover a resource fault condition.
 - OSCAGT_STATE_EVT
 - This OSC Event Class contains resource state change messages. Resource state change messages are sent by an OSC Agent when it detects a resource state change.

OSC Event Classes

- OSCSRV_CNXMAN_EVT
 - This OSC Event Class contains all OSC connection management messages from the OSC Service Engine.
 - The OSC Service Engine detects that the state of an OSC Agent has changed from operational to in-operational
 - The OSC Service Engine starts (restarts) OSC agents
 - The OSC Service Engine lost/established connection to the OSC Master Control Process
- OSCSRV_CONTROL_EVT
 - This OSC Event Class contains all control messages from the OSC Service Engine. The OSC Service Engine triggers a control message when it executes administrative operations on service groups and services at the request of the OSC Master Control Process.
- OSCSRV_STATE_EVT
 - This OSC Event Class contains all state change messages from the OSC Service Engine. A state change event is triggered if the state of one of the services and service groups managed by the OSC Service Engine has changed.

OSC Event Classes

- OSCCTRL_CNXMAN_EVT

- This OSC Event Class contains all OSC connection management messages from the OSC Master Control Process.

The OSC Master Control Process triggers connection management messages when:

- The OSC cluster state changes. These are:
 - » OSC cluster member leaves/joins the OSC cluster
 - » The OSC Master Control Process detects that the state of a managed OSC Service Engine has changed from operational to in-operational
- The OSC Master Control Process starts (restarts) an OSC Service Engine
- The OSC Master Control Process initiates an OSC cluster state transition
- The OSC Master Control Process modifies OSC quorum due to OSC cluster reformation

OSC Event Classes

– OSCCTRL_CONTROL_EVT

- This OSC Event Class contains all control messages from the OSC Master Control Process. The OSC Master Control Process triggers control messages when it executes administrative operations on service group (online, offline, switchover, failover, freezing a service group, disabling a service group ...) either due to a service group fault condition or on user request.

– OSCCTRL_STATE_EVT

- This OSC Event Class contains all state change messages from the OSC Master Control Process. A state change message is triggered when the state of a service group on any OSC cluster member changes.

OSC Event Severity

- Informational
 - Informational messages indicate that a managed item has changed state as expected or no state change related events have occurred. E.g.:
 - An informational event message is sent when a user connects to the console of the OSC Master Control Process via the OSC\$MGR utility.
 - An informational event message is sent if a user requests to set a service group offline and the service group changes state to offline as expected.
- Warning
 - Warning messages are triggered when OSC detects a fault condition that does not cause OSC intervention. Typically warning messages are sent if a non-critical resource fails or OSC rejects a user command due to state or authorization conflicts.
 - E.g. the user issues a command to set a service group online and the service group is already online, or the user is not authorized to manage a particular service group.

OSC Event Severity

- Error
 - Error messages are triggered when OSC detects any fault condition that causes OSC to intervene automatically. OSC is able to recover from the fault condition. Error messages are typically triggered when a resource, service and service group is declared faulted, a node is removed from the OSC cluster or OSC detects that one of its component (OSC Agent, OSC Service Engine, OSC Master Control Process) is in-operative and this component has had to be restarted.
- Fatal
 - Fatal messages are sent when OSC detects a fault condition that cannot be managed by OSC and requires immediate system management intervention. E.g. if a resource, service or service group state is set to ADMIN_WAIT a fatal event message is sent. A managed item is set to ADMIN_WAIT when OSC cannot recover the fault condition.

OSC Event Notification Methods

- Log events into the common OSC event message file
- Provide OPCOM messages for OSC events
- Forward events to all connected OSC consoles
- Execute a user script for each OSC event for site specific event handling

Bundled OSC Agents

- OscAgtDSK – OSC disk agent
- OscAgtFSYS – OSC single disk volume agent
- OscAgtSHD – OSC shadow set agent
- OscAgtPRC – OSC process agent
- OscAgtETH – OSC ethernet adapter agent
- OscAgtFailIP – OSC failSAFE IP agent
- OscAgtIP – OSC service IP agent
- OscAgtORA – OSC Oracle 10 agent
- OscAgtORALS – OSC Oracle 10 Listener agent
- OscAgtRDB – OSC Oracle RDB agent
- OscAgtPERF – OSC HP PERFDAT agent
- OscAgtMySQL – OSC MySQL agent
- OscAgtSQLSRV – OSC SQL service agent
- OscAgtSQLDIS – OSC SQL dispatcher agent
- OscAgtNCLObj – OSC DECnet object agent
- OscAgtDECnet – OSC DECnet alias agent

Supported Versions

- OpenVMS V7.3-2 AXP
- OpenVMS V8.2 AXP
- OpenVMS V8.3 AXP
- OpenVMS V8.4 AXP
- OpenVMS V8.2 IA64
- OpenVMS V8.2-1 IA64
- OpenVMS V8.3 IA64
- OpenVMS V8.3-1H1 IA64
- OpenVMS V8.4 AXP
- Any new OpenVMS versions of supported 8 weeks after shipment

HP OpenVMS ServiceControl distribution model

- OSC is Freeware
 - Free use for OpenVMS clusters up to 32 OpenVMS cluster members
 - Free use until 1-JAN-2100 (this is not a typo)
 - Any user of OSC agrees that he will share any changes and modifications (due to design og OSC possible for everyone) with all HP OpenVMS ServiceControl users.
- On customer request (not for free)
 - Customization and configuration support
 - Software support

HP OpenVMS ServiceControl links

- If you want more information about HP OpenVMS ServiceControl and/or if you want to be added to the HP OpenVMS ServiceControl update distribution list please contact:
 - Wolfgang Burger / HP Austria:
<mailto:Wolfgang.Burger@hp.com>
 - HP OpenVMS ServiceControl support link:
 - <mailto:ServiceControl@hp.com>
- Download Link:
 - <http://www.openservicecontrol.org/>

Q&A



Break



OSC configuration cook-book

- Plan your configuration according to the configuration rule
- Create a new OSC configuration database and configure all items according to your configuration plan
- Verify the configuration
- Define your newly created OSC configuration database as the default OSC configuration
- Test your configuration using the simulation mode of OSC
- If the test succeeds activate the configuration


OSC configuration rules

- The OSC cluster definition section of an OSC configuration database has to contain at least one member of the OpenVMS cluster you want to run OSC
- The OSC failover policy has to be defined in the OSC cluster definition section
- OSC service group, service and resource names have to be unique
- All OSC service groups have to contain a valid service group execution OSC node list
- An OSC service group has to contain 1..n OSC services
- An OSC service can be assigned to only one OSC service group


OSC configuration rules

- An On-Off resource can be assigned to only one OSC service
- On-Only and Persistent resources can be members of different OSC services and thus they can be members of different OSC service groups.
- On-Off resources can depend on other:
 - On-Off resources
 - On-Only resources
 - Persistent Resources
- On-Only resources can depend on other:
 - On-Only resources
 - Persistent Resources
- Persistent resources can depend on other:
 - Persistent Resources

OSC configuration Example




Node Failure




1. Define Service Groups

- 3 Service Groups in this example
 - HP-COM#1
 - Failover Service Group containing
 - HP\$GATEWAY_1
 - all resources this process depends on
 - HP-COM#4
 - Failover Service Group containing
 - HP\$GATEWAY_4
 - all resources this process depends on
 - HP-EXEC
 - Parallel Service Group containing
 - HP\$ROUTER
 - HP\$EZ\$AGT_1
 - HP\$EZ\$AGT_2 (VMSTM1 only)
 - HP\$SECMGR
 - all resources these processes depend on


HP-COM#1



HP-COM#4



HP-EXEC



Assign unique OSC resource names

Resource	OSC agent	OSC resource name	Remarks
HP\$ROUTER	OscAgtHERMES	HERMES::HP\$ROUTER	Node specific (VMTM4: no dependency on HP\$EZ\$AGT_2)
HP\$EZ\$AGT_1	OscAgtHERMES	HERMES::HP\$EZ\$AGT_1	---
HP\$EZ\$AGT_2	OscAgtHERMES	HERMES::HP\$EZ\$AGT_2	VMSTM1 only
HP\$SECMGR	OscAgtHERMES	HERMES::HP\$SECMGR	---
HP\$GATEWAY_11	OscAgtHERMES	HERMES::HP\$GATEWAY_11	
HP\$GATEWAY_41	OscAgtHERMES	HERMES::HP\$GATEWAY_41	

Assign unique OSC resource names

Resource	OSC agent	OSC resource name	Remarks
16.55.41.54	OscAgtFAILIP	FAILIP::HP\$GATEWAY_11	Node specific (IE0 on VMSTM1, WE0 on VMSTM4)
16.55.40.49	OscAgtFAILIP	FAILIP::HP\$GATEWAY_41	Node specific (IE0 on VMSTM1, WE0 on VMSTM4)
FAILSAFEIP	OscAgtPRC	PRC::FAILSAFEIP	On-Only – configure using /ADVANCED qualifier
\$1\$DGA93	OscAgtFSYS	FSYS::DGA93	---
\$1\$DGA99	OscAgtFSYS	FSYS::DGA99	---

Assign unique OSC resource names

Resource	OSC agent	OSC resource name	Remarks
DSA400	OscAgtSHD	SHD::DB1	Use AUTO/NOCFG_DISK to configure this resource
DSA410	OscAgtSHD	SHD::DB2	Use AUTO/NOCFG_DISK to configure this resource
DSA420	OscAgtSHD	SHD::DB3	Use AUTO/NOCFG_DISK to configure this resource
DSA430	OscAgtSHD	SHD::DB4	Use AUTO/NOCFG_DISK to configure this resource

OSC cluster settings

OSC cluster attribute	Value
OscCtrlClusterName	"Bootcamp-2010"
OscCtrlNode	VMSTM1:1, VMSTM4:1
OscCtrlFailoverPolicy	STATIC
OscCtrlProcPriority	[Unchanged default]
OscCtrlReconnInterval	[Unchanged default]
OscCtrlStartupWait	[Unchanged default]
OscCtrlStartupWait	[Unchanged default]
OscCtrlExpVote	1
OscCtrlAutoAdjustQuorum	Yes

OSC configuration

1. Create new OSC configuration project
2. Configure OSC cluster attributes
3. Create new Agent OscAgtHERMES
4. Create all resources according to the resource list (slides 68-70)
5. Create the Services and Service Group defined in the configuration plan (slides 65-67)
6. Define the resource dependencies as defined in the configuration plan
7. Verify and activate the OSC configuration project

Q&A

