



IBM Software Group

HAManager Overview and Common Issues

Kumaran Nathan(kumaran@us.ibm.com)

WebSphere Application Server Level 2 support

August 19 2014



WebSphere® Support Technical Exchange



Agenda

- High Availability (HA) Manager
- HA Services and Users
- HA Manager Configuration
- Core Group
 - ▶ Moving Nodeagent and Deployment Manager
 - ▶ Deleting a core Group
 - ▶ Core Group Policies
- HA Manager Resource Consumption
- Disabling HA Manager
- HA Manager Coordinator
- HA Manager Leader
- Tuning HA Manager
- TroubleShooting
- Common Issues

HA Vs Non-HA (Transaction recovery Scenario)

- Failure Scenario:
 - ▶ Client calls EJB™ which updates database using 2PC transactions
 - ▶ Failure during in-process transaction (after the prepare statement)
 - Database record is locked until transaction is recovered (committed or rolled back)
- V5
 - ▶ Server needs to be restarted and recover the transaction log
 - ▶ Could be up to 5 minutes
 - Other applications could be locked out if they required the same record
- V6 and Later
 - ▶ HA Manager detects the failure
 - ▶ Failover to a peer server which recovers the transaction log (shared on a NAS) from the failed server
 - ▶ Recovery in a few seconds

High Availability(HA) Manager

The HAManager Framework (HAM)

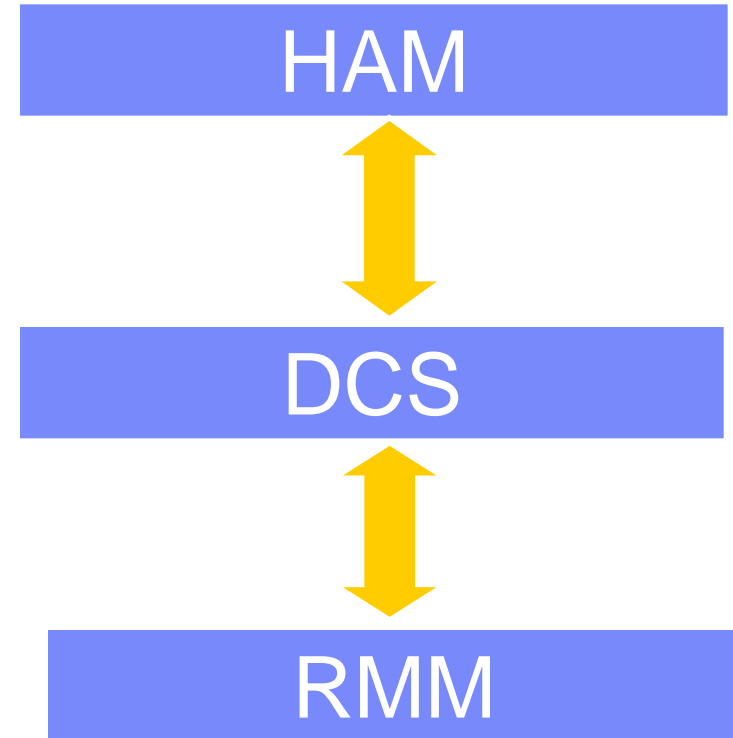
- ▶ An integral part starting WAS* 6.0 (and WAS XD) designed to provide an infrastructure for making selected WAS services highly available.
 - Present in all JVM™s including Deployment Manager and Node Agents
- ▶ Can be used by other internal WebSphere components to provide automatic failover support.
- ▶ Significant only in ND** environments
 - HAM provides no extra value for Base or any configuration consisting of a single server.
- ▶ Provides various asynchronous callbacks for interested internal WAS components.
- ▶ The Configuration of High Availability Systems is **Simplified**
 - Works out of the box in most cases - No additional administration required for most commonly used topology

*WAS – WebSphere Application Server

** ND – Network Deployment

HAManager Stacks

- Distribution and Consistency Services (DCS)
 - ▶ Default FIFO message ordering
 - Messages are delivered in the originating order
 - ▶ Optional virtual synchronous messaging
 - Either all or no group members receive a message
- Reliable Multicast Messaging (RMM) provides
 - ▶ TCP-IP like quality of service for multicast
 - ▶ Multicast emulation for TCP communication
 - Message fragmentation and reassembly
 - Congestion control
 - Transport layer for DCS
 - ▶ Tuning options
 - Transport buffer size, heart beat interval
 - Defaults are generally sufficient



HAManager Services

- Four basic HAManager services
 - ▶ **Bulletin Board (BB)**
 - A server state data exchange mechanism
 - Typically used for carrying routing information (**WLM and ODC**)
 - Scope: can be configured to span core groups (CGBS)
 - ▶ **High Availability Groups (HAGroups)**
 - Provides the HAManager policy based fail-over mechanism
 - Users : ME, Transaction Manager, JCA Adapter and more
 - The most visible service
 - Scope: Core Group
 - ▶ **Agent Framework**
 - Provides hot backup service from primary to N backup members
 - High throughput service used by **DRS** for data replication
 - Scope: Core Group
 - ▶ **Partitioned Managed Group (PMG)**
 - A distributed state manager (conceptually a distributed cache).
 - Used by the **core group bridge service** to forward bulletin board communication between core groups
 - Scope: Core Group

Present Users of HAM

- Messaging Engine
- Transaction Log Recovery
- All UCF (Unified Clustering Framework) based WLM
 - ▶ HTTP through WebSphere Proxy Server
 - ▶ SIP
 - ▶ EJB
 - ▶ Default Messaging
- Data Replication Service
 - ▶ HTTP Session
 - ▶ Dynacache
 - ▶ DistributedMap
 - ▶ Stateful Session Bean
- WebSphere Process Server
 - ▶ Highly Available Inbound JCA Resource Adapter
 - ▶ Event Sequencing
- WebSphere Extended Deployment
- Others??



HAManager Configuration

- Two levels of Configuration

- ▶ Cell scoped and Server scoped

- Cell Scoped (coregroup.xml file)

- ▶ `$WAS-PROFILE_ROOT\config\cells\<<cellname>\coregroups\<<core_group_name>\coregroup.xml`
- ▶ Servers > Core groups > Core group settings > CoreGroupName
- ▶ Contains Core Group attributes, policies, and members (WAS servers)
- ▶ **Changes in this file may not ALWAYS need a server restart**
 - Core Group level attribute changes are picked up dynamically
 - Custom Property introduction or changes should generally be followed by a restart of all the core group members

HAManager Configuration

■ Server Scoped (hamanagerservice.xml file)

- ▶ \$WAS-PROFILE_ROOT\config\cells\- ▶ Servers > Server Types > WebSphere Application servers > (server name) > Core group service (under additional properties)
- ▶ Mainly resource related
 - Thread pool. HAM shares threads with DCS
 - Server based HA Manager isAlive interval (default)
 - Default configuration should be good enough for most installations
 - Any alteration to this file requires a server restart

Core Group

Core Groups

- Core Group (HA Domain)
 - ▶ Logical grouping of application servers.
 - ▶ A set of Application Server processes over which direct High Availability relationships may be established.
 - ▶ Core Group information is persisted in a configuration file.
 - ▶ WAS install creates the *DefaultCoreGroup*, which can never be deleted.
 - ▶ Server creation/deletion adds/removes entries from the corresponding core group file.

Core Groups - Formation Rules

- Core groups must conform to the following rules
 - ▶ A core group may contain zero or more WebSphere processes (JVM's)
 - ▶ A core group containing enabled members must contain at least one administrative process (until 6.1.0.x)
 - deployment manager
 - node agent
 - ▶ A WebSphere process (JVM) must belong to one and only one core group
 - ▶ A core group cannot span cells.
 - All the members of a core group must belong to same cell
 - A cell can contain more than one core group
 - ▶ A cluster cannot span core groups
 - All of the cluster members must belong to the same core group
 - A core group can contain more than one cluster



Core Groups

- Typically, the single *DefaultCoreGroup* definition will suffice
- In a large WAS topology additional core groups may be required.
 - ▶ Scalability
 - ▶ Performance
 - ▶ Existence of Firewall between cell members
 - ▶ Geographically separated cell members
 - ▶ Also a placeholder for HAManager disabled servers
- A core group is an equal opportunity container
 - ▶ No special status for node agents or deployment manager.
- CRUD operations on Core Groups can be performed from ISC* or using wsadmin scripts.

*ISC – *Integrated Solutions Console*



Moving A Node Agent and DMgr

- For moving a node agent across core groups
 - ▶ Stop the node agent
 - ▶ Perform the movement to the desired core group
 - ▶ Save and sync the configuration changes to the running nodeagents
 - ▶ Perform a `syncNode` operation on the node agent that was moved
 - ▶ Start the node agent
- If you must move the deployment manager from core group CG1 to core group CG2, perform the following steps
 1. Stop the deployment manager process
 2. Start a “local mode” scripting session under the dmgr profile
 - `wsadmin -conntype NONE -lang jython`
 3. In the local mode scripting session move the deployment manager from CG1 to CG2
 - `AdminTask.moveServerToCoreGroup("-source CG1 -target CG2 -nodeName kumaranNode -serverName dmgr")`
 4. Save the configuration changes
 - `AdminConfig.save()`
 5. Restart the deployment manager
 6. Synchronize the configuration changes out to the running nodes.



Deleting a Core Group

- A non-empty core group can not be deleted

- For deleting a core group
 - ▶ Move all the servers to other core groups maintaining the core group formation rules
 - ▶ Delete the empty core group



Core Group Policies

- Policies determine the fail-over characteristics of the WAS components using the HAGroup service.
- By default two policies are automatically created in the DefaultCoreGroup.
 - ▶ Default SIBus Policy – for Messaging Engines, and WPS Event Sequencing
 - ▶ Clustered TM Policy – for Transaction Manager
- Policy names must be unique in a *coregroup.xml* file
- CRUD operations on policies do not require a server restart
- Do not modify or delete the two default policies
 - ▶ If necessary create your own policies



Core Group Policies

- Five types of pre-defined Core Group policies
 - ▶ **OneOfN** Policy
 - HAM will activate the component on one server in the HA Group.
 - Most popular policy.
 - ▶ **Static** Policy
 - HAM will activate the component on statically configured servers in the HA Group.
 - Manual intervention needed for fail over.
 - ▶ **NoOp** Policy
 - No HAM initiated activation/deactivation callbacks to components.
 - ▶ **MOfN** Policy
 - HAM will activate the component on M ($M \geq 1$) servers in the HA Group.
 - Not in use
 - ▶ **AllActive** Policy
 - HAM will activate the component on all the servers in the HA Group.
 - Not in use.

HA Manager Resource Consumption



HA Manager Resource Consumption

- HA Manager component uses resources for Discovery, Failure Detection, and View Synchrony.
- Up to 50 member core groups, resource usage is minimal
 - ▶ CPU
 - Idle per process CPU usage is typically minimal (< 1 %).
 - CPU spikes can occur when view changes, but are short lived.
 - CPU utilization concerns are typically related to z/OS customers who are charged per MIP.
 - ▶ Network Bandwidth
 - Network congestion due to Failure Detection (heart beat) is not an issue on a standard gigabit network
 - Heart beat is optimized and not sent if messages are already flowing
- Recommendation: start with 50(or 100) members per core group.
 - If CPU/Memory adequate can be larger but do not exceed 100.

HA Protocols

- **The View Synchrony Protocol**

- ▶ Responsible for establishing reliable messaging with other core group members after the connections are opened.

- **The Discovery Protocol**

- ▶ Responsible for discovering when other core group processes start, and for opening network connections to these other members.

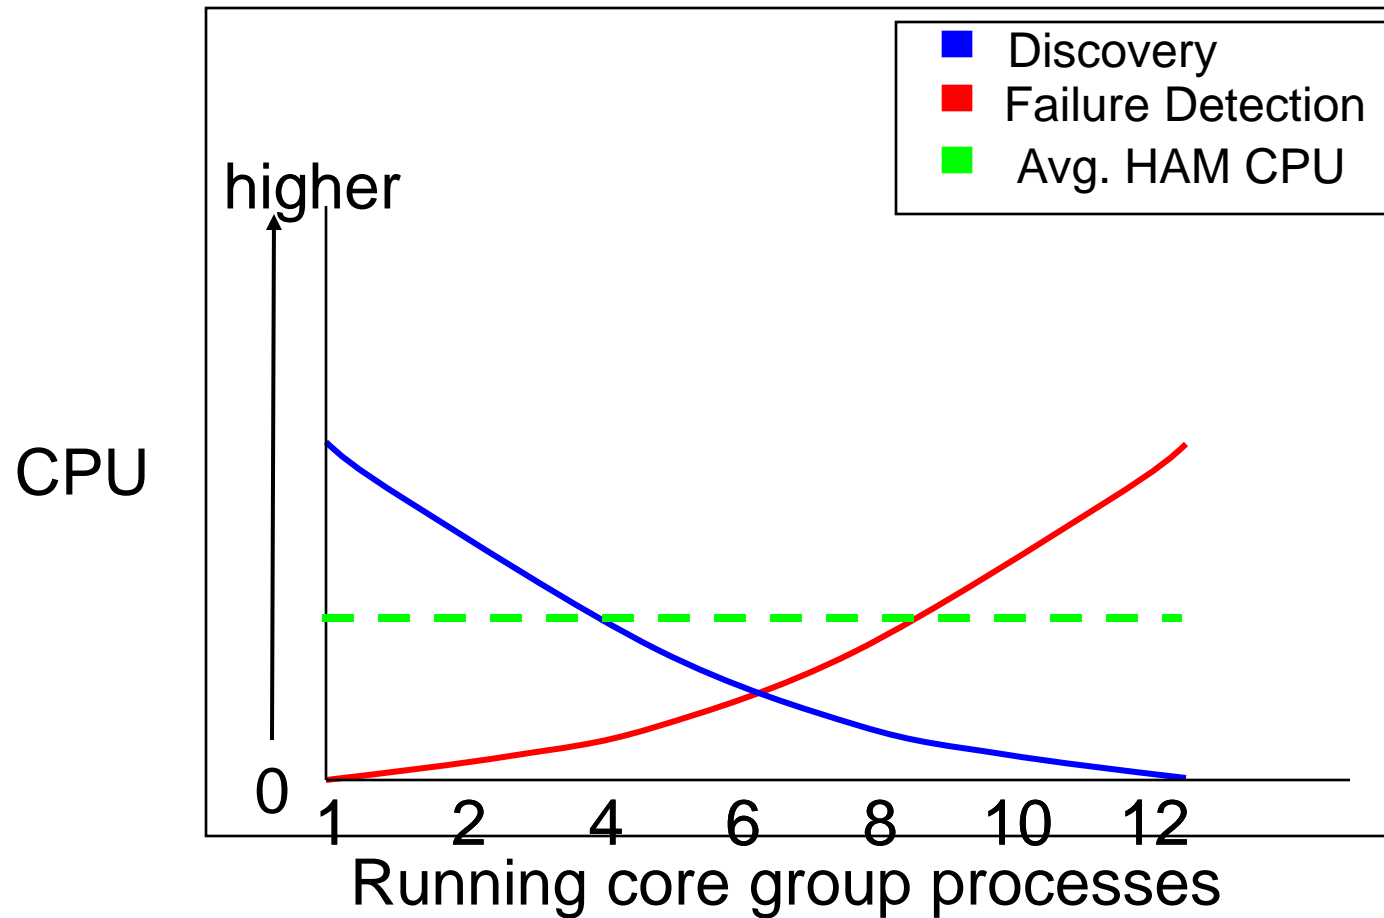
- **The Failure Detection Protocol**

- ▶ Responsible for detecting when other core group members stop or become unreachable because of a network partition.

- These protocols are executed in all processes.
- If required the protocols can be customized or tuned.
- In most topologies, the default settings are sufficient and no tuning is required.



HA Manager Resource Consumption - conceptual



Disabling HA Manager

Disabling HAManager

- **HAManager and the components** that use it's services consume resource.
 - Memory
 - CPU
 - Network Bandwidth

- It is possible to disable the HA Manager services if they are not used.
- Disabling HA Manager will effectively disable all the components using HA Manager services.
- Every release of WebSphere has new components that leverage the HA Manager services.
 - ▶ This makes it difficult to determine the HA Manager dependencies.



Disabling HAManager

- HA Manager services are required for any of the following:
 - ME fail over in a cluster
 - Transaction Log Recovery in a cluster
 - Replication Domain members
 - EJB's (RMI-IIOP) in a cluster
 - Proxy Server (SIP, HTTP, WS-AT, WS-Addressing)
 - Remote Session Invalidation (RRD)
 - WebSphere XD
 - WebSphere Process Server (Inbound JCA RA, Event Sequencing)
 - etc.
- **RECOMMENDATION**
 - ▶ Do not disable the HA Manager service unless you are absolutely sure the service is not being used now or will not be used in future.

Disabling HAManager

- Recommended approach for disabling HAManager
 - ▶ Create a new core group, say *DisabledCoreGroup*
 - ▶ Move all the relevant servers, for which you want to disable HAManager into the *DisabledCoreGroup*.
 - ▶ Write a script to disable HAManager in all the members of *DisabledCoreGroup*.
 - The *DisabledCoreGroup* can contain many JVMs
 - No limits on the number of JVMs in *DisabledCoreGroup*
 - No extra resources get consumed
 - There is no need to have a managed process (node agent or deployment manager) in *DisabledCoreGroup*.

Disabling HAManager

- Starting WAS 6.1, HAManager can be disabled using ISC by
 - Unchecking the “Enable service at server startup” check box, followed by
 - Server restart

Cell=kumaranaCell01, Profile=Dmgr01

Application servers

[Application servers](#) > [server1](#) > Core group service

Use this page to specify core group related settings for the application server.

Configuration

General Properties

Enable service at server startup

* Core group name

DefaultCoreGroup

Disabling HAManager

■ Effect of HAManager disablement

- ▶ Identical to highly optimized HAManager in a WebSphere Base or Express server.
- ▶ (Optimized) HAManager classes with no DCS or RMM
- ▶ No heart beat checking at DCS level.
- ▶ No cross process messages or bulletin board traffic
- ▶ Resources for HAM and components using HAM released
- ▶ For a HAManager disabled server, every restart may log HMGR0005 informational message in the SystemOut.log.
 - [4/14/14 17:51:54:141 CDT] 0000000a DCSPluginSing I HMGR0005I: The Single Server DCS Core Stack transport has been started for core group. <<Disabled Core Group message>>

HAManager Coordinator

- The (active) **coordinator** aggregates distributed state information from the individual processes.
 - ▶ By default there is only one active coordinator per core group.
 - ▶ It's role is to manage the location of the services that depend on the HAManager (HAM) for high availability.
 - ▶ By default, HAManager selects the lexicographically lowest named server from the available core group members.
 - ▶ You can configure a preferred coordinator server(s)
 - An **ordered** list of servers can be specified
 - Servers > Core groups > Core Group Settings > <Core Group Name> > Preferred coordinator servers
- The active coordinator, should be located on a **stable** server that is not restarted often.
 - ▶ Sufficient memory – for holding the state and other HAGroup information
 - ▶ Light load – for quickly reacting to the core group membership change events. The Deployment Manager, Node Agents, or stand alone application server can be considered for hosting coordinators.
- A single active coordinator does not introduce a single point of failure

View Leader

- The view leader is responsible for coordinating activity between the core group members and then checking with the **active coordinator(s)** to confirm the change with the HA stack.
- The view leader role is not user configurable and is chosen by WebSphere.
- The view leader makes the internal changes necessary to reflect the event and confirms the new set with the active coordinator(s) for the core group. If everything checks out OK, the new view is installed and all core group members are notified of the change. WebSphere emits a log entry when a view change occurs :
- ***DCSV8050I: DCS Stack DefaultCoreGroup at Member Kumarancell\testnode\server1: New view installed, identifier (1090:0.Kumarancell\dmgrnode\dmgr), view size is 26 (AV=26, CD=26, CN=26, DF=40)***
- The message tells us several key pieces of information:
 - ▶ 1. A view change occurred in the DefaultCoreGroup core group
 - ▶ 2. A new view was installed,
 - ▶ 3. The view serial id number is 1090 meaning there have been 1089 previous incarnations of the view for this core group,
 - ▶ 4. The view leader for this core group is currently the Dmgr server and the view size is 26 servers.

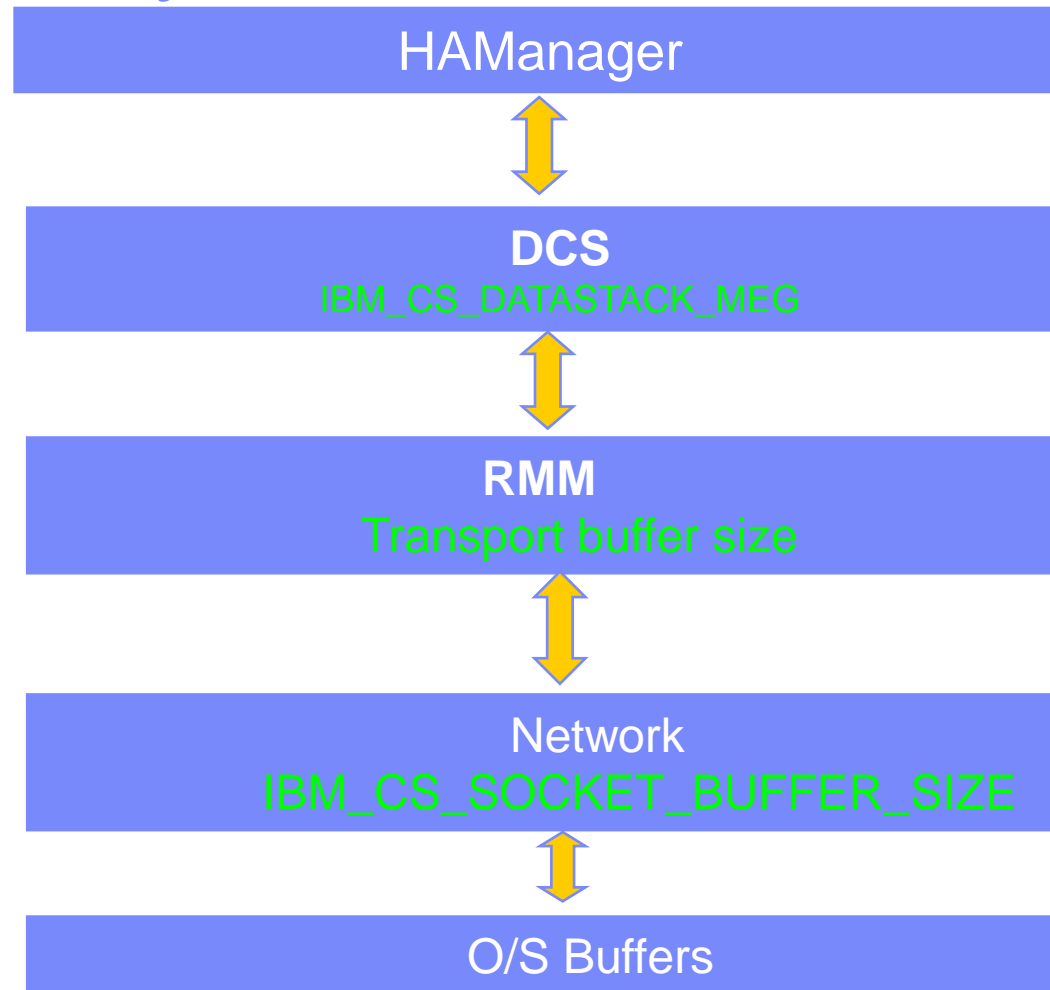
HA Tuning , Troubleshooting and Common Issues



Tuning HAM – Memory Buffers

V6.1 and earlier release only

- Recommended if using Replication or Large Topology
 - ▶ change IBM_CS_DATASTACK_MEG, a core group custom property to 100
 - ▶ change the server-scoped 'Transport buffer size' attribute to 100 for all the core group members.
 - ▶ restart all the members
 - ▶ increase the value of these two parameters, if DCS congestion messages continuously appear in SystemOut.log



Tuning HAM – Memory Buffers Cont...

- The high availability manager consumes its allocated heap memory when it sends interprocess messages on behalf of other services that use the high availability manager functionality. For example, the heap memory might be consumed while sending memory-to-memory replication data, or highly available routing data from one core group member to another core group member.
- In Version 7.0, the new transport memory size setting at **Servers => Core Groups => Core group settings => CORE_GROUP_NAME** replaces the `IBM_CS_DATASTACK_MEG` custom property and the per process transport buffer size setting. The transport memory size has a default value of 100 MB, so setting this property is no longer necessary.
- It is not required to set `IBM_CS_SOCKET_BUFFER_SIZE` custom property. The default value is sufficient in most of the cases. Visit [Knowledge center](#) for more information.

Transport buffer size

- In large topologies default value of 100 MB is adequate. The following messages may be appearing in the SystemOut.log which indicate problems transferring large amounts of message data.

DCSV1051W, a high severity congestion event for outgoing messages

DCSV1052W, a medium severity congestion event for outgoing messages

DCSV1054W, a medium severity congestion event for incoming messages

- If you see any of the above message increase the Transport memory size value.
 - ▶ Default 100 MB
 - ▶ Maximum can't be more than JVM heap size

General Properties

* Name

Description

Default Core Group. The default core group cannot be deleted.

* Number of coordinators

* Transport memory size

megabytes

IBM_CS_OOM_ACTION

- Use this custom property to explicitly enable exception handlers that are specific for `OutOfMemoryExceptions` that occur when sending or receiving network messages. When this property is set to `Isolate`, if an `OutOfMemoryException` occurs when a network message is being sent or received, these exception handlers stop High Availability Manager communications to the Out of Memory process.
- `CWRLS0030W` may be avoided when the server fails to start due to OOM.
- Add the custom property under Core Groups > Core Group Name > Custom properties

Core Stack protocols

- **IBM_CS_WIRE_FORMAT_VERSION:** This property controls **the communication protocol** used by members of a core group. For large topologies, it is recommended to use the latest version that the topology can support since new efficiencies are introduced with each version.
 - ▶ Only 50 servers are recommended per coregroup with the default protocol version.
 - ▶ 100 servers are recommended per coregroup with the new protocol version.
 - ▶ Available protocol versions are 6.0.0, 6.0.2.9 and 6.1.0 but recommended version is 6.1.
- **IBM_CS_HAM_PROTOCOL_VERSION:** This property enables core group functionality that **improves core group bridge scalability**. It was added via APAR PK64003 to WebSphere Application Server service levels 6.0.2.31, 6.1.0.19, and 7.0.0.1. Regardless of version, the only value for enabling this property is "6.0.2.31."
- No server restart required when you add these custom properties.



HAManager Troubleshooting

- Important HMGRxxxx informational messages
 - ▶ **HMGR0228I**: The Coordinator is not an Active Coordinator for core group DefaultCoreGroup.
 - ▶ **HMGR0206I**: The Coordinator is an Active Coordinator for core group DefaultCoreGroup.
 - ▶ **HMGR0207I**: The Coordinator was previously an Active Coordinator for core group DefaultAccessPointGroup but has lost leadership.
 - ▶ **HMGR0218I**: A new core group view has been installed. The core group is DefaultCoreGroup. The view identifier is (9:0.<view_leader>). The number of members in the new view is 1.
 - ▶ **HMGR0005I**: The Single Server DCS Core Stack transport has been started for core group DefaultCoreGroup.
 - ▶ **HMGR0226I**: The core stack configuration parameter {0} has been set to {1}.

HAManager Troubleshooting

- Important HMGRxxxx warning/error messages
 - ▶ **HMGR0024W**: An error was encountered while looking up the IP address for the host name of a core group member. The host name is {0} and the exception is {1}.
 - Error occurred resolving IP address for host name
 - ▶ **HMGR0027W**: A duplicate host to IP Address mapping has been detected. IP Address {0} has a mapping to host names {1} and {2}.
 - Duplicate host name to IP address mapping.
 - ▶ **HMGR0028E**: A duplicate DCS_UNICAST_ADDRESS port has been detected. Members {0} and {1} on host {2} are both configured to use port {3}.
 - Duplicate port.
 - ▶ **HMGR0152W**: CPU Starvation detected. Current thread scheduling delay is 9 seconds.
 - This could occur because CPU utilization is 100%, but it could also occur for other reasons (e.g. memory is exhausted and disk paging or garbage collection is causing scheduling delays). Can be ignored, if its not frequent and not more than 20 seconds.

HAManager Troubleshooting

- Important DCSVxxxx informational messages
 - ▶ **DCSV1032I**: DCS Stack DefaultCoreGroup at Member <this member>: Connected a defined member <remote member>
 - Indicates a connection to another core group member was established
 - ▶ **DCSV8050I**: DCS Stack DefaultCoreGroup at Member <this member>: New view installed, identifier (9:0.<view leader name>), view size is 27 AV=27, CD=27, CN=27, DF=27)
 - Indicates that a new view was installed.
 - AV (Actual View) – the number of members in the new view
 - CD (Connected - Denied) – the number of connected members minus the number of denied members (members are denied from the view for various reasons)
 - CN (Connected) – the number of connected members.
 - DF (Defined) – the total number of core group members defined in the configuration (but not necessarily running)
 - Typically, when things have settled, the AV, CD, and CN numbers will all be the same. DF will be different unless every member of the core group is running.
 - ▶ **DCSV8030I**: DCS Stack DefaultCoregroup at Member <this member>: Failed a join attempt with member <remote member>. The reason is Not all candidates are connected ConnectedSetMissing= [] ConnectedSetAdditional = []
 - Indicates a member failed to join the view. Can happen and be ignored if only occur intermittently. However if the message is logged continuously it is a sign of a communication problem.

HAManager Troubleshooting

■ Important DCSVxxxx warning messages

- ▶ **DCSV1111W:** DCS Stack DefaultCoreGroup at Member <this member>: Suspected another member because the outgoing connection from the other member was closed. Suspected members is <remote member>. DCS logical channel is View|Gossip.
- ▶ **DCSV1112W:** DCS Stack DefaultCoreGroup at Member <this member>: Suspected member <remote member> because of heartbeat timeout. Configured Timeout is 180000 milliseconds. DCS logical channel is Connected|Ptp.
- ▶ **DCSV1113W:** DCS Stack DefaultCoreGroup at Member <this member>: Suspected another member because the outgoing connection to the other member was closed. Suspected member is <remote member>. DCS logical channel is Connected|Ptp.
- ▶ **DCSV1115W:** DCS Stack DefaultCoreGroup at Member <this member>: Member <remote member> connection was closed. Member will be removed from view. DCS connection status is Discovery|Ptp, transmitter closed.
 - Are all indications that a socket closed. These can be expected when a process is administratively stopped or killed. These messages indicate the remote process that closed the socket will be removed from view. If these messages occur, but the remote process was not stopped or kill, it is an indication of some other problem.
- ▶ **DCSV8104W:** DCS Stack DefaultCoreGroup at Member <this member>: Suspected member [<remote member>] because the member is suspected by member <remote member>. Internal details VL suspects others: CC-Situation Normal
 - Indicates that a process will be removed from view, based on information from another view member. In other words, this process did not yet detect any problems with the remote member, but another process in the view is instructing this process to remove the remote member from view.

HAManager Troubleshooting

■ Important DCSVxxxx warning messages

- ▶ **DCSV1134W**: DCS Stack DefaultCoreGroup at Member <this member>: Suspected another member upon receiving an additional connection from the other member. Suspected member is <remote member>. DCS logical channel is Discovery|Ptp.
 - A 2nd (additional) connection was received from a member. Often this is the result of duplicate ports in the configuration where two members are attempting to open connections using the same port.
- ▶ **DCSV1051W**: DCS Stack {0} at Member {1}: Raised a high severity congestion event for outgoing messages. Internal details are {2}.
- ▶ **DCSV1052W**: DCS Stack {0} at Member {1}: Raised a medium severity congestion event for outgoing messages. Internal details are {2}.
 - Indicate congestion is occurring. Should be OK if the message is only occasionally logged. If the messages are regularly logged the message buffers sizes should be increased.
- ▶ **DCSV1053I**: DCS Stack {0} at Member {1}: Outgoing messages congestion state is back to normal.

Common Issues

- Network configuration issue
 - ▶ Multiple NIC, Host/IP Resolution
- Configuration issue
 - ▶ Corrupted server.xml file, serverindex.xml file, coregroup.xml file
- Network issue
 - ▶ Firewall, Network Partition (can cause split view)
- Server start issue
 - ▶ Server fails to start with CWRLS0030W
 - ▶ Refer the [Technote](#)



Additional HA Resources

- **Automate peer recovery for transactions and messages in WebSphere Application Server V6.0.x**
http://www.ibm.com/developerworks/websphere/techjournal/0509_lee/0509_lee.html
- **Transactional high availability and deployment considerations in WebSphere Application Server V6**
http://www.ibm.com/developerworks/websphere/techjournal/0504_beaven/0504_beaven.html
- **CWRLS0030W message continuously logged and WebSphere Application Server fails to open for e-business**
<http://www.ibm.com/support/docview.wss?uid=swg21245012>
- **Coregroup Bridge**
http://www.ibm.com/developerworks/websphere/library/techarticles/0606_petersonk/0606_petersonk.html
- **CWRLS0030W message continuously logged**
<http://www.ibm.com/support/docview.wss?uid=swg21245012>
- **DCSVXXXX message in SystemOut.log file**
<http://www.ibm.com/support/docview.wss?uid=swg21188934>

Summary

- High Level Overview of HA Manager
- Importance of HA Manager
- Advantage and Disadvantage of disabling HA manager
- HA Manager Tuning Tips
- HA Manager Troubleshooting

Connect with us!

1. Get notified on upcoming webcasts

Send an e-mail to wsehelp@us.ibm.com with subject line “wste subscribe” to get a list of mailing lists and to subscribe

2. Tell us what you want to learn

Send us suggestions for future topics or improvements about our webcasts to wsehelp@us.ibm.com



Questions and Answers

Additional WebSphere Product Resources

- Learn about upcoming WebSphere Support Technical Exchange webcasts, and access previously recorded presentations at:
http://www.ibm.com/software/websphere/support/supp_tech.html
- Discover the latest trends in WebSphere Technology and implementation, participate in technically-focused briefings, webcasts and podcasts at:
<http://www.ibm.com/developerworks/websphere/community/>
- Join the Global WebSphere Community:
<http://www.websphereusergroup.org>
- Access key product show-me demos and tutorials by visiting IBM Education Assistant:
<http://www.ibm.com/software/info/education/assistant>
- View a webcast replay with step-by-step instructions for using the Service Request (SR) tool for submitting problems electronically:
<http://www.ibm.com/software/websphere/support/d2w.html>
- Sign up to receive weekly technical My Notifications emails:
<http://www.ibm.com/software/support/einfo.html>