



Control-M

Application Continuity Version 9.0.21.100

Table of Contents

Introduction	3
Document Purpose	3
What is Application Continuity?	3
Why is Application Continuity Important?	4
Control-M as a High Availability Application	4
Document Contents	6
Hardware Solutions.....	7
Disk Storage Replication	8
Control-M Implementation with Disk Storage Replication	8
Host and Virtual Machine.....	9
Virtual OS clustering Operating System Clustering	9
General Information about Clustering	10
Failover Scenario in Clusters	11
Application Layer	13
Database Solutions.....	13
Database Replication	13
Standby Databases.....	13
Cluster Database Solutions	16
Guidelines for Control-M installation with Parallel Server Cluster Solution.....	16
Control-M Highly Available (HA) Configuration	16
Control-M/Enterprise Manager High Availability Architecture using an Oracle, MSSQL and external PG database	17
Control-M/Server High Availability Architecture	18
Control-M high availability on one host.....	19
Control-M/Enterprise Manager distributed instances.....	20
Control-M/EM and Control-M/Server high availability with Oracle/MSSQL/ePG	20
Control-M/Server high availability with BMC PostgreSQL.....	21
Control-M/Server High Availability Architecture with PostgreSQL	22
Network Latency	23
Control-M Replication Using File Copy.....	23
Proactive Notification Tools	24
Frameworks Integration	24
Best Practices for Full Control-M Databases Backups.....	24
Best Practice for Performing Backups with BMC Provided PG in Control-M HA Mode	27
Best Practices for Disaster Recovery	28

For More Information	32
Where to Get the Latest Product Information	33

Introduction

Document Purpose

Application Continuity in the production environment is a critical issue that every modern organization faces. This document aims to provide information about Application Continuity and why it is important. It also reviews the available solutions for achieving Application Continuity in Control-M, as well as guidelines for their implementation.

The target audience for this document is software consultants, technical support staff, and customers. Since it is imperative that this document is kept as updated and as accurate as possible, any information or remarks are welcomed and can be sent directly to “CustomerSupport@bmc.com”.

This document includes examples of configuration and implementation of Application Continuity solutions. These examples are based on practical experience gathered in the lab and the field. Not all of the examples were tested according to the official standards, and these examples should therefore be treated as recommendations only, with limited liability.

What is Application Continuity?

Application Continuity is defined as the means taken by an organization to ensure that critical applications, such as Control-M, function continuously under all circumstances. When a failure is localized and limited to an application in a given location, we refer to the term **High Availability** as a means to ensure Application Continuity. For Control-M, High Availability (HA) is defined as the means taken to minimize the frequency and duration of downtime in the production environment in a specific site with no data loss to scheduling and/or Workload Automation. Downtime can be caused by either unexpected hardware or software failures, or by planned events, such as scheduled maintenance like OS version updates. The target of any organization that wants to establish High Availability levels is to maintain the lowest possible number of unexpected downtime incidents and the shortest possible planned or unplanned downtime.

With HA solutions in place, organizations can ensure Application Continuity against any disaster that occurs within a given location. However, if a disaster completely destroys a primary site, then the organization must have a contingency plan that replicates the primary site with a secondary site some distance away. The

planning for a contingency site in a situation where the primary site has been completely disabled will be referred to as planning for **Disaster Recovery (DR)**.

Organizations do not always require both HA and DR solutions. If the distance between primary and secondary sites is within the technical limitation for High Availability solutions, then organizations can reduce the expense of building both HA and DR within their organization. Technology has improved to a point where the ability to build an HA environment across hundreds of miles is more cost effective for most organizations.

Why is Application Continuity Important?

The importance of Application Continuity lies in the profound effect it has on modern business success. Modern regulatory and availability requirements make the need for Application Continuity an integral part of today's IT Infrastructure. The trend towards globalization (and the internet in particular) and cloud services has brought availability issues to the forefront — even for industries that did not pay much attention to it in the past. Many organizations that not so long ago defined “acceptable downtime” as a few hours, today aim to reduce downtime to minutes or even seconds. Recurring and lengthy downtime has a direct effect on critical business aspects, including:

- Revenues are lost during downtime.
- Employee productivity is reduced.
- Customer satisfaction and confidence is sometimes lost.
- Technical support expenses increase.

Consequently, today's organizations are investing in solutions that ensure their applications' availability. However, implementing and maintaining High Availability is a major expense and must be included in the cost of the applications and services that IT provides. Usually captured in Service Level Agreements (SLAs), cost justification includes weighing the potential loss resulting from downtime and impact on any regulatory or audit requirements. Thus, each organization must determine the required level of availability for each of their applications and associated services and select the option best suited to their needs.

Control-M as a High Availability Application

Control-M is a best-of-breed workload automation solution that addresses business needs by focusing on applications, integration, and management of Workload Automation. It integrates automated workflows and batch processes running on diverse platforms and applications into a unified business process, ensuring service management requirements and smooth operations. Since Control-M triggers critical application jobs, its availability level must be extremely high — any downtime may affect major business services that involve

numerous other critical applications. *Figure 1* (page 5) presents the basic architecture of the Control-M product, which is comprised of several components that can each be installed on a different machine.

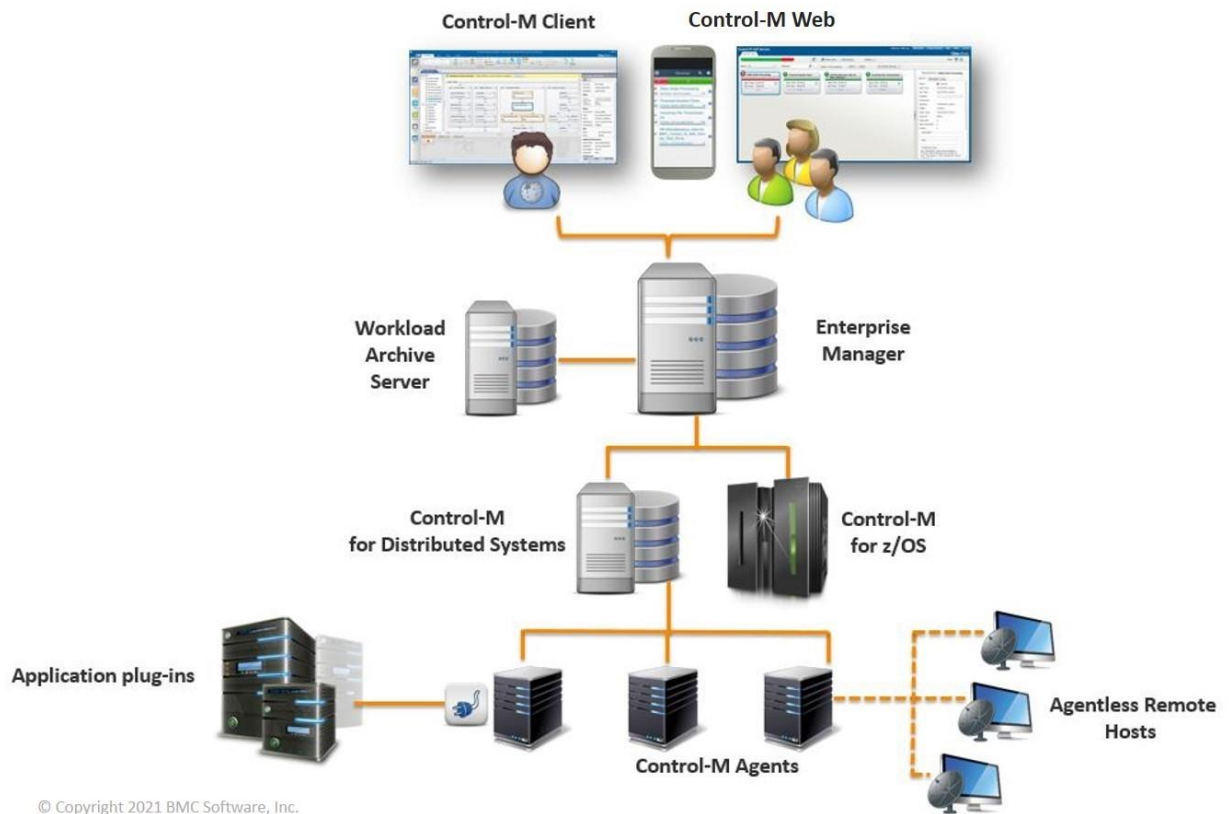


Figure 1: Basic Control-M Architecture

This flexible and scalable architecture is designed in a way that if one component on one machine fails, or the communication fails, the other components in the system continue to provide service based on the most recent available data, when possible.

In this document, Application Continuity refers to all possible means available to keep Control-M working on its vital task — Workload Automation and job scheduling. From a hardware perspective, Control-M makes extensive use of databases and, as such, is vulnerable to disk failures. Database redundancy should be included in the Application Continuity strategy. Application Continuity can be approached from several angles, and each angle can be implemented autonomously. In most cases, the integration of these efforts will maximize

the availability of the system. Every organization should choose the configuration that best suits its environment, needs, and resources.

Document Contents

This document reviews the available solutions in each of the following areas, and discusses and reviews the integration of each solution with Control-M:

- **Hardware Solutions:** Reviews common available fault tolerance systems, which overcome CPU, memory, network, and disk failures.
- **Operating System (OS):** Discusses Virtual machine and OS management, including OS clustering solutions for all supported platforms, which minimize the downtime after hardware or software failures.
- **Database Solutions:** Reviews database vendor solutions that are currently available in the market for maintaining database stability and their possible integration with Control-M.
- **Control-M HA Configuration** – Discusses the built-in Control-M High Availability solution and its integration with Control-M.
- **Proactive Notification Tools:** Presents the proactive notification solutions available from BMC and other vendors, which integrate with Control-M either to enable prevention of downtime or to rapidly identify a failure so that manual intervention can take place.
- **Best practices for Backing up Control-M:** Discusses the recommended way to back up the database irrespective of what solution you implement for High Availability or Disaster Recovery.

The following diagram shows the layers of a typical application “stack” in today’s modern datacenter. This model uses four layers and shows the abstracted components in each layer.




Figure 2: Typical Application Stack

From a Control-M perspective, the main consideration is that the types and effects of the layers below the application and database are virtualized and abstracted from the actual hardware. This is by design, and it provides additional levels of availability. The physical or environmental requirements for the application and database are mapped to the virtual machine, instead of the physical hardware.

An often-overlooked initial consideration is the need for bi-directional communication between Control-M components in both the hardware (physical) and the virtual network.

Hardware Solutions

The traditional fault tolerance model relies on specialized hardware that detects a hardware fault and instantly switches to a redundant hardware component — whether the failed component is a processor, memory board, power supply, I/O subsystem, or storage subsystem.

Disk Storage Replication

Disk replication is a fault tolerance solution that is transparent to the user. It provides a reliable replication of the environment that can be used in specific failover scenarios. There are many vendors that provide real-time data replication, such as EMC2 with SRDF to Hitachi True Copy, DRDB for Linux®, and IBM® PPRC. These vendors support different operating systems and different levels of High Availability storage, such as RAID 1+0, RAID 0+1, and RAID 1=0+0. High Availability storage elevates the availability of applications that use it.

Control-M Implementation with Disk Storage Replication

Data replication can be used to replicate the entire Control-M environment from one machine to the other only in situations where the secondary machine assumes identically all machine attributes. These attributes include, but are not limited to:

- Platform and OS level must be identical between primary and secondary machines.
- Kernel parameters must be identical.
- OS user definition, group, and its internal IDs must be identical.
- Home directory of the OS user and its environment must be identical.
- Hostname should be the same (or see [Knowledge Base](#) for [Control-M/Enterprise Manager](#) and [Control-M/Server](#) for instructions on what changes are required when renaming the hostname).

Data replication, when it meets the requirements listed above, is supported for both HA and DR configurations. The contents of the Control-M account can be replicated for either the application files only, or it can be replicated for both application files and database files. The database file replication must follow the database vendor's guidelines and supported practices for real-time data replication. Specific database vendor documentation should be consulted on instructions on how to implement real-time disk replication. It is very important to follow guidelines specified by the database vendor — especially since data files used by the database server are constantly updated and the data consistency must be adhered to (keeping in mind that Control-M's consistency is based on a committed transaction). In most cases, when the database file replication is set to synchronized mode, the data consistency is protected.

Because the guidelines for real-time data replication (RDR) vary greatly between the Control-M application and third-party databases, when using RDR for meeting failover demands, the application and database components should be documented and maintained separately in the overall strategy.

Keep in mind that Control-M/Enterprise Manager database is accessed only when information is needed for most activities except from the Gateway. The Gateway is similar to online transaction processing and requires a

high throughput capability. On the other hand, Control-M/Server mostly requires online transaction processing capabilities for managing continued job processing. Control-M/Server connects to its own database schema and is independent of the Control-M/Enterprise Manager database.

Host and Virtual Machine¹

Virtualization technology provides a layer of abstraction between the computing, storage, and networking hardware, and the software that runs on it. This technology enables users to run additional operating systems in multiple sessions via a Hypervisor called virtual machines (VMs).

BMC welcomes the running of its products in the leading virtual environments. The VM environment can be configured to provide additional levels of availability independently from the Control-M configuration. In general, BMC products fully support running under virtual environments. For additional information, see the Virtualization policy on our web site.

<http://www.bmc.com/support/policy-for-virtualized-platforms.html>

Virtualization in a cloud environment, such as Amazon, Azure, or Google, adheres to the same policy as above. However when dealing with SaaS/PaaS databases, such as RDS, the support for the database services is not provided unless it is documented in the Compatibility section for each component. An RDS solution provides redundancy and ensures database continuity that fits in the overall schema of application continuity. Selecting a database that is external to Control-M, such as an RDS option, also facilitates additional solutions with Control-M high availability that are not possible when selecting an internal database option.

Virtual OS clustering Operating System Clustering

Virtual OS clustering involves clustering multiple virtual machines (VMs) on a single physical server, with the goal of providing high availability and fault tolerance. In this setup, multiple VMs are created and distributed across multiple physical servers, with software and hardware mechanisms in place to manage and monitor the virtualized environment. Virtual OS clustering is often used in cloud computing environments, where the virtual machines can be quickly migrated to other physical servers if a server failure occurs.

Operating system clustering, on the other hand, involves clustering multiple physical servers to work together as a single, highly available system. In this setup, multiple servers are configured to work together in a cluster, with

shared storage and networking resources, and software that enables them to function as a single system. The primary goal of operating system clustering is to provide high availability and fault tolerance for critical applications and services, and to minimize downtime in the event of hardware or software failures.

Both virtual OS clustering and operating system clustering are important techniques for building highly available and fault-tolerant IT infrastructures. However, they differ in their scope and the level of abstraction they operate at. Virtual OS clustering is concerned with managing the virtualized environment and ensuring the availability of virtual machines, while operating system clustering is concerned with managing physical servers and ensuring the availability of critical applications and services.

General Information about Clustering

Cluster computing architecture is used in many organizations to address expanding IT needs and requirements for High Availability. This architecture is based on deploying groups of computing resources, including CPUs, memory, disks, and network interfaces. Cluster computing addresses these current IT issues by providing increased service availability, increased scalability, better resource management, and improved manageability.

A cluster consists of two or more machines connected together with specialized hardware and software that work together to form a seamless single machine. Disk storage is usually external and accessible to both machines, although that is not necessarily the case. Each machine is referred to as a node. Machines do not have their own unique users; adding users to the cluster makes the user login available everywhere. Neither does each machine require its own user licenses; these are shared between the nodes. Each node has its own IP address, and the cluster itself has a virtual IP address. In most cases, the servers in the cluster are interconnected using a faster network interface, such as Fiber Channel connection, for better performance. This network is called the Heartbeat Network.

The cluster software monitors the condition of the nodes, and whenever it identifies a failure in CPU, memory, network, disk, or application, it addresses another node to provide the required service. In addition to the failover feature, clustering provides for workload balancing across the computer. *(Please note that this feature is not discussed in this document).*

The cluster management software can vary between simple, custom, in-house programs and large products offered by the major vendors. The hardware can range from a PC to a Parallel Sysplex mainframe.




Figure 3: Basic Cluster Architecture

Failover Scenario in Clusters

The most common High Availability solution in cluster environments is a **hot standby server**. A hot standby server is a host that is dedicated to providing all of the services of another host. The host that is the primary provider of the services is called the primary server. Every server that is part of a cluster is called a node. In case of a problem in the primary server, the hot standby server takes over control and provides the services of the primary server. The process of taking another node's services is called failover or takeover.

The failover process always starts with the detection and identification phase of the failure condition. The failure may occur at the hardware level, at the OS level, or even at the application level. An example of a failure can be a memory problem, system panic, or a database server crash. The detection and identification can be performed manually or automatically. Console automation products and frameworks can perform the automatic detection internally using cluster management software or externally using outboard automation products.

One of the most typical examples of a failover process is the recovery scenario from a complete cluster node crash. The heartbeat monitor of a console automation product detects the lack of response by one of the cluster's nodes. The console automation verifies that it is not just a network timeout, but also a hardware failure on this specific node, and starts the failover process.

The system administrator defines failover scenarios for each failure condition, using the vendor-provided cluster management software package. When the failure condition is identified, its corresponding failover scenario is triggered. Failover triggering can be performed manually by the operator, internally using the cluster built-in mechanisms, or by outboard automation products.

In this typical scenario of a complete node crash, the designated backup node receives a command that triggers the failover process. As a result, the following steps occur:

1. The backup node turns on one of its additional network interfaces and assigns the failed node IP address to it. This enables the backup to respond to client service requests, which are addressed to the failed node IP.
2. All of the previously available services and applications are started on the backup node.

In most cases, this transition is transparent to the end users and client applications, which typically experience the transition as an application timeout. In some cases, however, users may be required to restart their connections.

Furthermore, in most cases, the manual recovery process is performed during the least critical time period for the business.

Control-M/Server cluster support is limited to Active/Passive cluster configuration. For a typical Control-M failure scenario, the production resumes from the point where it stopped. The installation directory and the database reside on the shared drive and are accessible to the standby server. The database contains the most recent state of production at the moment of failure, and the standby server serves as a substitute for the primary server in all aspects. All jobs that were executing on the host that encountered a failure are aborted. Jobs that were running on the agents are not affected at all by the server's failure; they continue to run and their connection to Control-M/Server is reestablished when the standby machine has reinitialized the Control-M application.

Each of the Control-M components can be configured as a separate cluster application group. The entire application group is treated as an entity, and if any resource in the application group fails, the entire application group is taken offline and restarted on the second node. In a Windows platform, Control-M integrates with Windows Cluster API to install an application and configure it as an out-of-the-box clustered application. For Unix platforms, Control-M provides a ready script required for any Unix cluster to manage any standard application. When Control-M runs in a cluster, the EM components cannot be managed by CMS and instead the Cluster manages the components through Control-M provided scripts.

For Control-M implementation and guidelines for installing in UNIX, Linux, or Windows environments, see the Control-M Installation Guide chapter on "[Cluster Configuration](#)".

Application Layer

Application layer redundancy is an important technique for ensuring the high availability and reliability of critical application components, providing a better user experience, and minimizing downtime due to hardware or software failures.

Database Solutions

The purpose of database redundancy is to ensure high availability and data recovery in the event of a hardware or software failure. Control-M customer data is all resident in the database. Having the ability to recover application data allows the user to have the ability to standup a new environment and restore the working data.

Database Replication

Database Replication is a software solution that enables the user to define tables and parts of tables to be replicated to another database. The term primary is used to describe the instance of the database that is being updated. Replication can either be Multi-Master, where each database can be updated, or Primary Snapshot, where there is one primary database and the changes are propagated to read-only snapshot sites referred to as replica. In addition, the replication can be set up in two modes:

- **Synchronous mode:** Each transaction ends only after distributing it to the entire replicated database.
- **Asynchronous mode:** The updates to the replicated sites are distributed at certain intervals or upon request.

Synchronous mode is not suitable when the primary and replica have a network latency greater than five milliseconds (traceroute). Examples of vendor-specific database replication services are Oracle Data Guard, GoldenGate, PostgreSQL replication, and MSSQL replication.

Replication servers secure the consistency of the application data only. Database binary files, log files, and other configuration data owned by the database server are not replicated with a given replication tools.

Standby Databases

Oracle® Standby Server

Oracle Standby Server (called Data Guard in Oracle 12c and higher) is a hot standby solution for the database server. This solution is software-based and is considered to be a less costly option than cluster solutions. In addition, there is no geographical limitation regarding the location of the standby database, so it can be used as part of a Disaster Recovery plan. On the other hand, to support High Availability there is a need to enable

synchronous mode of database replication. Additionally, there are guidelines for maximum network latency for a replicating server working in a synchronous mode. In a DR scenario where asynchronous replication mode can be enabled, the network latency considerations can be relaxed.

The configuration consists of two database servers; one is the primary database server and the other is in a standby mode. The servers can either run on different machines linked via a regular communication line, or they can run on different machines in the same subnet. The primary server distributes its log files to the standby machine that can apply the changes immediately or not, based on the configuration. The detection of a failure and the failover procedure is a series of steps that can be operated manually or contained in a script that can be run automatically by the system that monitors the database. After failover, which causes some inevitable downtime, the standby database will function as the main database server. Returning to the original database server is not a one-step process.

Installation of Control-M with Oracle Standby Server

No special guidelines are required for installation with Oracle Standby Server. However, it should be noted that in case of failover and a switch to the standby database, Control-M should be stopped and reactivated again, in order to reconnect to the database.

In a High Availability scenario, the Control-M application can either:

- Continue processing by repointing Control-M environment to the standby database, or
- Manually start a standby Control-M environment that includes a preinstalled Control-M application that points to the standby database server already.

For specific details on how to *configure* Control-M to point to Oracle Standby Server, refer to the description of the 'restore_host_config' utility in the Control-M Utilities Guide.

In asynchronous mode, database replication is recommended for Control-M/Server in a Disaster Recovery scenario only. Control-M/Enterprise Manager in asynchronous mode is suitable for both HA and DR scenarios. In a DR scenario, a full version of Control-M is required to be installed on the secondary site, including HA configuration with similar configuration as the primary site. The reason why asynchronous replication is suitable for Control-M/Enterprise Manager in an HA scenario is because the active environment can be forced to resynchronize when the Control-M/Enterprise Manager comes up for the first time in a standby configuration and reconnects to the Control-M/Server. *To force synchronization for each data center, see the Troubleshooting menu of **root_menu** or see the documentation for the Control-M/Server utility **reset_ecs**.*

Copying Control-M Oracle Database

When using the Oracle-provided utilities to copy an Oracle database from a primary environment to a replica database server, it is important that the user definitions are not modified in the target database server. Oracle User roles should remain, and you should not move user individual privileges to a group role. Stored procedures in Oracle behave differently when a privilege is set at individual level or at the role level. `Restore_host_config` script should be used to localize the data in the database. Running `restore_host_config -dr` script also ensures that the configuration settings used in the primary are replicated from the database to the configuration files of the secondary. Refer to **ctmfilesync** utility for details.

Installation of Control-M in Oracle Replication Database Environment

Database replication cannot be defined for the Control-M database, since this requires certain database schema attributes that are not currently part of the product database schema.

PostgreSQL Replication

When utilizing Control-M standalone or the Control-M HA solution, PostgreSQL server can also provide asynchronous database replication using hot backup with archiving enabled. The archiving creates an archive log at configured intervals or based on size limitation that is partitioned to a new archived file. This replication process requires an automated script in place to ship the archived files over to the standby Control-M environment and either restore the shipped files immediately in the standby database server or leave them to be restored in case of disaster. This PostgreSQL solution is suitable for a disaster recovery scenario because replication can be done only in asynchronous mode.

One can utilize both Control-M HA and archived enabled hotbackup at the same time to maintain both HA and DR environments. PostgreSQL manages the two features independently. Restoring a hotbackup on a machine with a different hostname requires you to localize the data first, before it is restored. Refer to KA [000323891](#) for details.

When replication is done in a DR environment where DR Control-M is installed in a standby mode and the database is in read only mode, you must have a full plan when applying Fixpacks or patches in the DR environment. Upgrading to a higher release or applying a Fixpack and patches requires additional steps, because the standby database is in read-only mode. For instructions, refer to the Installation Guide section "Upgrading Control-M in a Disaster Recovery configuration". Setting environment variable `inst_skip_db_changes` to Y allows the upgrade to take place without performing any of the SQL commands to update the read-only database. This option is available starting with Control-M 9.0.19.

Cluster Database Solutions

Oracle, MSSQL, and PostgreSQL RDS offer solutions that use the benefits of the cluster environment to enable activation of multiple SQL servers over the cluster nodes. These solutions (Oracle RAC Parallel Server, Oracle AWS RDS, MSSQL AlwaysON, Amazon Aurora, and PostgreSQL RDS on AWS and Azure) provide transparency during the failover process by taking over the database session of the failed node in a way that is totally transparent to the application. The communication between the database servers on each node is very fast due to the FDDI communication between the cluster nodes, each node having its own database memory structures, and the database files residing on the shared disk of the cluster on several raw devices. A cluster database solution is preferred over disk replication because in addition to storage replication it also provides database server redundancy.

With the 9.0.19 release of Control-M support, Control-M now supports several RDS databases in AWS cloud, Azure cloud, and Google cloud. An RDS solution can provide a similar level of database continuity as a clustered solution. Refer to your provided cloud for details on different levels of redundancy that are offered in cluster solutions.

Guidelines for Control-M installation with Parallel Server Cluster Solution

Control-M fully supports installation with Oracle RAC server. There are no differences when installing Control-M with a standard database server or with RAC server. The installation is transparent to Control-M — as long as you use Service Instance Name and not SID Name.

Control-M Highly Available (HA) Configuration

Control-M High Availability enables you to maximize your production environment uptime and prevent data loss in the event of hardware maintenance or failure. Control-M supports the following high availability solutions:

- Control-M/EM and Control-M/Server high availability with Oracle, MSSQL, or external PostgreSQL enable you to set up a secondary host with Control-M/EM or Control-M/Server. If there is a hardware failure or if all Control-M/EM or Control-M/Server processes are down unexpectedly, the secondary automatically (by default) or manually assumes control and resumes production.
- Control-M/Server high availability with internal PostgreSQL: Enables you to set up a secondary Control-M/Server and a secondary PostgreSQL database server for database replication. If the primary Control-M/Server and database server are down, you can manually fail over to the secondary host. This option does not exist for Control-M/Enterprise Manager.

Control-M High Availability is a software solution that monitors if the entire application is up or down. It is not a hardware level solution that monitors all aspects of hardware and conditional levels of the application health.

To set up your high availability environment please see the Administration and Installation guide found on https://documents.bmc.com/supportu/9.0.21/en-US/Documentation/Introduction_to_Control-M_installation.htm

Control-M Workload Automation client will automatically detect the active EM server and route the connections to the appropriate EM web server component.

Control-M/Server with PostgreSQL server replication replaces the Control-M/Server mirroring provided in previous versions.

Control-M/Enterprise Manager High Availability Architecture using an Oracle, MSSQL and external PG database

The following diagrams show the configurations of both the Control-M Enterprise Manager (EM) and Control-M Server for high availability using Oracle, e/PG, or MSSQL databases.

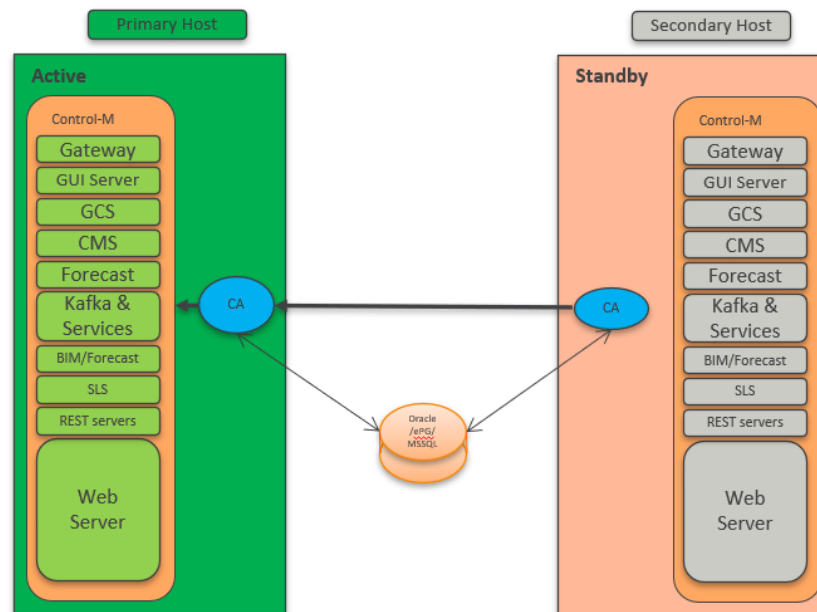


Figure 4: Control-M/EM with Oracle/MSSQL/external PG

The following diagram shows the configuration of the Control-M/EM after automatic failover when the primary components are no longer available. High Availability is provided only on the Primary EM instance, the distributed EM instances are not included in high availability solution.

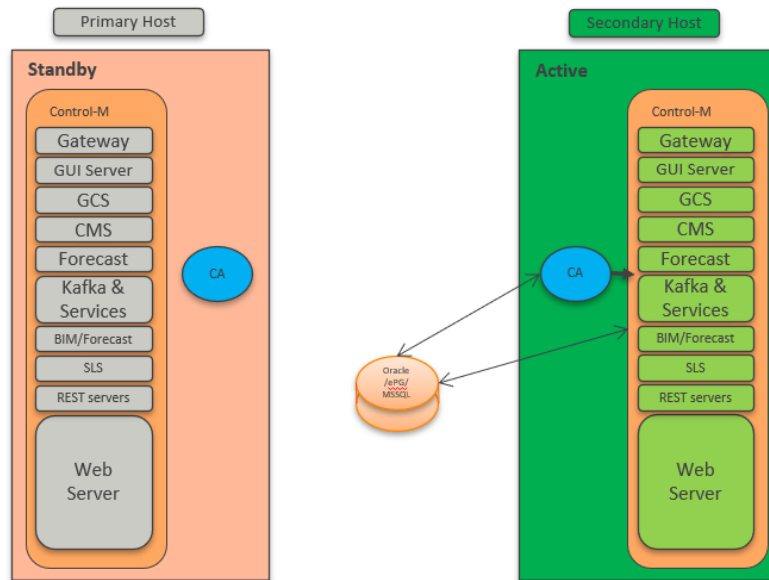


Figure 5: Control-M/EM High Availability Configuration

Control-M/Server High Availability Architecture

The following diagram shows the high-availability configuration of the Control-M/Server using an Oracle, external PG, or MSSQL database.




Figure 6: Control-M/Server High Availability Configuration

The following diagram shows the configuration of the Control-M/Server using an Oracle, external PG, or MSSQL database after automatic failover when the primary components are no longer available.

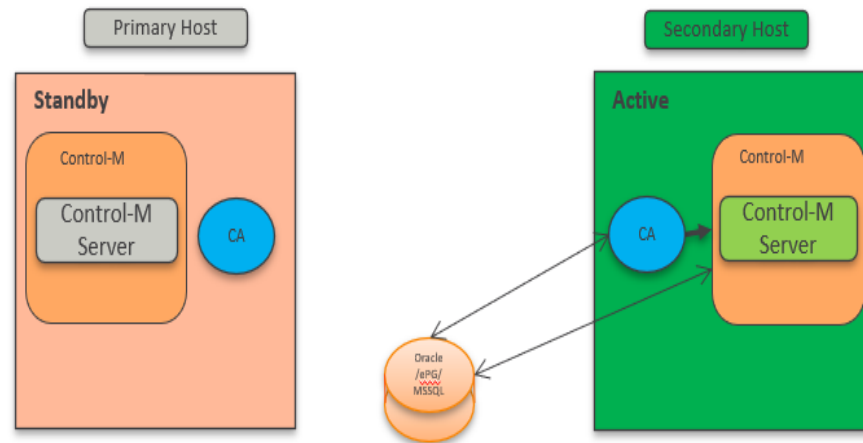


Figure 7: Control-M Server Automatic Failover

Control-M high availability on one host

To ensure that Control-M continues to run when it is installed on one host, the Control-M/EM and Control-M/Server Configuration Agent monitor and manage the following components in the CCM:

- Control-M/Server
 - Control-M/Server
 - Services (Kafka, Zookeeper, Scheduling, etc.)
 - PostgreSQL database server
- Control-M/Enterprise Manager
 - GUI Server (GSR)
 - Gateway (GTW)
 - Global Condition Server (GCS)
 - Batch Impact Manager Server (BIM)
 - Forecast Server
 - Self Service Server

-
- Web Server
 - Configuration Manager Server (CMS)
 - Services Configuration Manager
 - PostgreSQL database server
 - Services (Kafka, Zookeeper, Reports, etc.)

If a Control-M/EM or Control-M/Server component goes down, the Configuration Agent, if configured for automatic restart, will attempt to start it up (if desired state is set to Up) based on defined intervals, as described in Maintenance parameters and High Availability parameters.

This configuration of One Install is supported only with Oracle, external PG, and MSSQL databases.

Control-M/Enterprise Manager distributed instances

A distributed instance of the Enterprise Manager (EM) installation represents a single running copy of the EM application. It functions as an independent execution environment separate from the Primary/Secondary EM instance. All EM instances connect to each other to form a unified EM installation with a single database schema. Distributed EM instances are specifically required for Workload Archiving and Workflow Insights installations.

It's important to note that a distributed EM instance is not part of Control-M High Availability and does not provide redundancy. However, starting with Control-M 9.0.21, Workload Archiving supports an Active-Active configuration, allowing the installation of multiple distributed EM instances with Workload Archiving. This setup ensures redundancy of Workload Archiving in case a single distributed EM instance fails.

With multiple Workload Archiving instances, each instance shares the load of archiving job details. In the event of a failure, the workload is distributed among the remaining Workload Archiving instances. A similar architecture is supported for Workflow Insights as well.

Control-M/EM and Control-M/Server high availability with Oracle/MSSQL/ePG

After you have installed a secondary Control-M, the Control-M/EM, Control-M/Server, Configuration Agent on the secondary host monitors the primary Control-M/EM or Control-M/Server based on defined intervals. The monitoring is performed on two aspects. The first is with heartbeat over the network between the primary Configuration Agent and the secondary Configuration Agent. The secondary CA sends a life check to the primary CA and waits for a confirmation. The second aspect is that the secondary CA monitors a table in the

database that both share for a current timestamp of the last transaction. When both aspects of the test fail, you can fail over to the secondary in one of the following modes:

- Automatic failover: The secondary Configuration Agent automatically takes control and resumes production, when it detects that the primary Control-M/EM or Control-M/Server and its primary Configuration Agent has stopped unexpectedly.
- Manual failover: You can perform a manual failover at any time from the CCM if the manual failover option is enabled. After the failover is complete, the production runs on the secondary.

NOTE: If you attempt to manually start up components on the secondary when the primary is active, the components shut down automatically. This prevents both the primary and secondary from running components simultaneously.

Control-M/Server high availability with BMC PostgreSQL

The Control-M/Server high availability solution with a BMC-provided PostgreSQL database supports both a manual failover of the Control-M server along with PostgreSQL database replication. Control-M/Server is installed on the same host as the PostgreSQL database on the primary and on the secondary.

After the data replication is turned on and initialized successfully, the Control-M/Server database data is replicated synchronously to the secondary database server. However, if there are network communication problems, the replication mode switches to Asynchronous. The flat file configuration data is replicated as well to a shared drive, which is used if the primary or secondary are temporarily unavailable. See the Administration Guide for prerequisites for the shared drive.

The secondary Configuration Agent monitors the primary to verify life check responses from Control-M/Server and that the primary Configuration Agent is working, based on defined intervals. You can perform a manual failover at any time from the CCM if the manual failover option is enabled, based on the conditions described in [Manual failover](#).

If you are using a BMC PostgreSQL database, the Configuration Agent manages the database component and sends a life check every defined interval (see [Maintenance parameters](#) and [High Availability parameters](#)). If there is no response after a defined number of attempts, the Configuration Agent restarts the database automatically. The restarting of PG is performed in the same environment. When the Database alone is down, High Availability is not triggered in this scenario. Only when Control-M processes are down, the High Availability functions become activated.

If you are using an Oracle, external PG, or MSSQL database, you can view the database component in the CCM, but the Configuration Agent does not manage the component and cannot start it up or shut it down.

The following diagrams show the Control-M Server high availability architecture using the PostgreSQL database.

Control-M/Server High Availability Architecture with PostgreSQL

The following diagram shows Control-M/Server in a high availability environment using a BMC-provided PostgreSQL database. BMC-provided PostgreSQL is different from an external PostgreSQL that is provided and supported by a customer. The ownership of an external PostgreSQL belongs to the customer. The advantage of an external PG is that HA configuration is supported by Enterprise Manager.

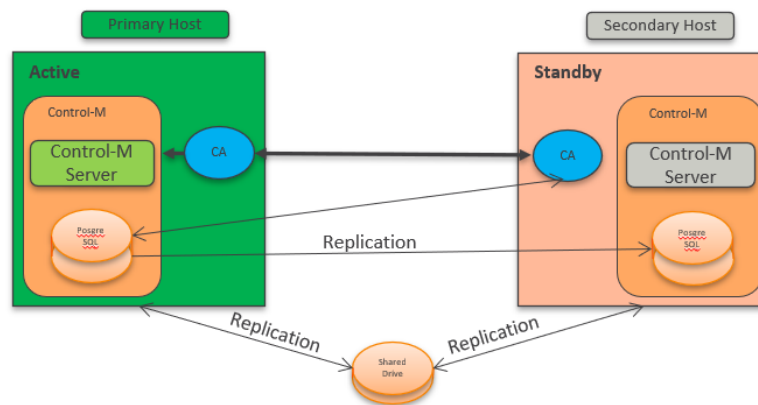


Figure 8: Control-M/Server HA with PostgreSQL

The following diagram shows a Control-M/Server manual failover when the primary components are no longer available.

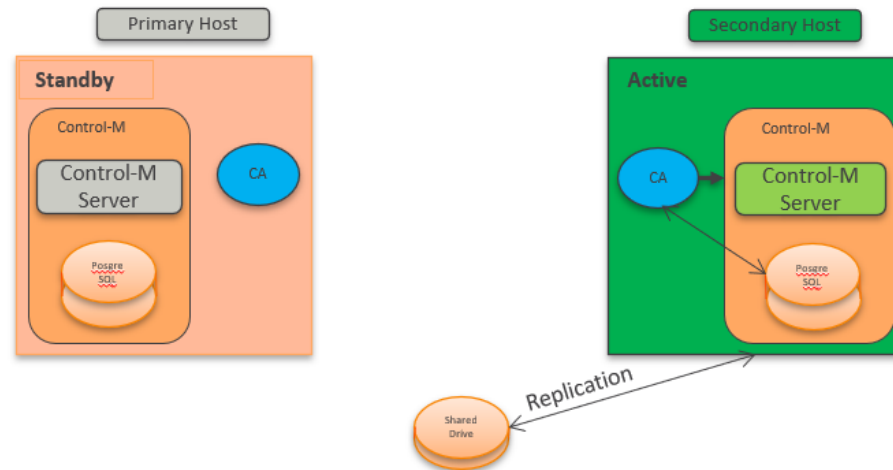


Figure 9: Control-M/Server with PostgreSQL Manual Failover

Control-M/Enterprise Manager High Availability with PostgreSQL database replication is not available with the v9 GA release.

Network Latency

The Network is the conduit that connects the primary host to the secondary host, the conduit that connects the application to the Database, and connects the Database to the replicated Database. In Control-M HA configuration we assume that that all components are located on-site and in a DR situation the components are separated across city or state lines where the network latency is measured in tens or hundreds of milliseconds. But even in the same location, where you would expect to see network latency in single digits, there are many factors that can impact the throughput, such as number of routers, firewalls, cable quality, etc. For best practices and guidelines on what is required to properly facilitate application continuity, refer to Knowledge Article [000231726](#)

Control-M Replication Using File Copy

Control-M currently **does not support** any form of file copy from a running and configured environment to a DR site or a standby secondary site. The only way to replicate an existing environment is to perform a full install of the secondary environment. The full install guarantees that all local configurations that are specified to the local environment are set properly. Control-M at installation time sets internal variables that point to the full path of Control-M environment, hostname, user account, and other local variables that cannot simply be reconfigured when an existing account is copied over to a new location.

If the database resided on the local machine, file copy of open files on a file database server is not supported by any database vendor. Each specific database vendor's documentation should be consulted for an acceptable means of how to copy the database files. Control-M includes database backup and restore utilities to facilitate an acceptable way of copying database data files. *See the Control-M Utilities Guide for more details.*

Proactive Notification Tools

This section reviews the tools that can be installed and integrated with different Control-M components in order to give proactive notifications about Control-M events. These notifications enable the user to prevent, or at the very least identify failures of the product components so that preventative or recovery actions (such as the activation of the failover mechanism) take place as soon as possible.

Frameworks Integration

The integration of Control-M with Frameworks allows better management and higher availability of Control-M by monitoring its different components.

This integration allows Frameworks to monitor different components of Control-M. Problems are detected immediately as they occur and then escalated to the Framework application, where corrective action can be taken to keep continuous production flow. Control-M provides ability to send SNMP traps to any network management framework such as HP OpenView or IBM Tivoli software. Control-M/Enterprise Manager can send application Alerts and administrative XAlerts via SNMP.

Additional information can be found in [Knowledge Article 000294783](#)

Best Practices for Full Control-M Databases Backups

To protect Control-M against database failures, backup of the database should be performed in such a way that one can recover to the point of failure. Control-M provides two ways to perform backups of the application database. The traditional way (or the **cold backup**) can be performed when the application is down. Cold backup is performed when no updates are being made to the database, thus protecting the integrity of the data. Control-M provides a second mechanism for backing up a live database without having to shut down the Control-M application. This is referred to as a **hot backup**. The integrity of the data is maintained in a hot backup because together with the running database snapshot, a hot backup includes a list of running transactions (more commonly referred to as the transaction log) that bridges the gap between the base

snapshot and the last committed transaction. The database snapshot together with the transaction log combine to recreate the database with full integrity.

A cold backup can be inconvenient to a mission-critical environment that needs to be up 24x7. Additionally, a cold backup will restore the data to the same point when the backup was performed. For Control-M/Server database this is not an ideal point of recovery. Control-M/Server database should be restored to the point of failure to ensure that job flow processing continues from point of failure and not from a few hours back.

Hot backup technology is dependent on a database feature called **archiving mode**. When archiving is enabled, the database incremental changes are written to the transaction log file, and then, after a configured amount of time or a configured size of transaction log, the log file is closed and made available for shipping to a safe designated storage. The transaction log files are sequentially labeled from the time the archiving is enabled. Every time a new hot backup is performed the archived logs are marked and can easily be identified as belonging a new backup snapshot image. The archived logs file should be stored on a safe storage location that is protected from machine and disk crashes.

When enabling archiving mode with a PostgreSQL database, the user is prompted for an empty folder where the archived log files will be written into when the transaction log files are closed. The active transaction logs file (also known as **Write Ahead Log**, or **WAL**, files) are kept in the default location `<pghome>/data/pg_xlog/`. They serve as the source files that feed the archive log files. The WAL file is managed by PostgreSQL and remains open at constant size for a specified duration. Only after the WAL file is closed does it get moved automatically to the archive log directory. Assigning a network drive for the archived log directory will ensure that files are safe from damage in case the host machine crashes.

Because the open WAL contains up-to-the-minute transactions, in case of a disk crash, it is in jeopardy of becoming damaged or corrupted. The transaction log files are simple flat files, and therefore, there is a risk that some of the data that was already committed is still buffered in the OS memory and not written out to the file. To manage the OS file buffering, a file synchronization mechanism exists with most database vendors. This solution ensures that for every database commit, the memory buffers are flushed, and the data is written to the disk.

The PostgreSQL parameter **fsync**, when turned on, ensures that for every database commit, the memory buffers are flushed, and the data is written safely to the disk. There is some performance penalty when the **fsync** parameter is turned to on. On databases with light activity, it is also recommended to add the **archive_timeout** parameter to the `postgresql.conf` file with a value of 300. This parameter forces an archiving activity every five minutes. Even if the WAL files are not fully used, it forces a WAL file to close and be shipped.

This also ensures that, in case of a disaster, the worst possible loss will be the past activities that occurred only in the last five minutes.

As the transaction log files fill up and are moved to the archive log directory, the number of archived log files can grow dramatically when there is a lot of activity. Therefore, in order to make the recovery process manageable, we recommend that a hot backup be made every 24 to 48 hours. The hot backup creates a new baseline, and thus makes all the previous archived log files irrelevant so they can be deleted. As soon as a hot backup is made, the baseline for a full restore is created and this baseline backup should be moved over to a safe location. Following the hot backup, all archived log files that will be created are for the incremental changes and will be now associated with the new hot backup baseline. Hot backup is initiated while Control-M is running.

For PostgreSQL database, when you initiate a new hot backup, you are prompted for an empty directory to store the backup files. The backup files will include both a baseline database file and a full set of database binary files. The database data file is an export of the entire database instance. Therefore, if both Enterprise Manager and Control-M/Server are serviced by the same instance, both tablespaces will be included in this export file. *Additional information can be found in [Knowledge Article 000311604](#).*

The following is a summary of the backup strategy discussed above

- 1. Enable archiving mode**

- This option is necessary in order to run hot backups.

- 2. Perform a hot backup**

- Create a baseline of the database by performing a hot backup. A hot backup can be executed any time without bringing down the Control-M application.
- Create a new baseline every 24 to 48 hours, depending on the amount of activity and the size of the archive log files.
- Purge logs archived prior to the hot backup after the new baseline is created.
- Ship the baseline backup file to a safe storage location immediately.

- 3. Ship archive log files**

- As soon as the transaction log file is closed, ship it to a safe archived log location.

Both hot and cold backup operations perform full database backup using database vendor utilities. They also create a backup that is application independent. In order to restore this backup, Control-M needs to be down, and the entire system is restored to the point of the last archived log file.

When running a hot restore in a primary or failover environment, ensure that archiving mode is off and that Control-M application is down. With PostgreSQL, you are also required to ensure that the PostgreSQL database is down prior to executing a hot restore. The hot restore utility will prompt you for both the archive location of the archived log files and for the location of the hot backup baseline. The hot restore brings the database to the last consistent position when the transaction log file was closed. The PostgreSQL restore also restores all the binary and database files to the point of the last backup. *For additional limitations on restoring a hot backup file with PostgreSQL, see Knowledge Database articles [000264656](#) and [000323891](#).*

For PostgreSQL database, the restore operation will import both tablespaces automatically when one instance is used to service both Control-M/Server and Enterprise Manager components.

A full database backup for Enterprise Manager can also serve for safekeeping previous versions of application data. Performing a backup every 24 or 48 hours allows you to restore the entire database to a previous date and time. In Control-M Enterprise Manager, this can be useful when a user mistakenly modifies application data, such as job definitions, and needs to restore to some point in the past. With a PostgreSQL database in situations where both Control-M/Server and Enterprise Manager reside in the same instance, the restore is not very practical. In order to restore application data from a full backup, you need to build or have ready a test or a failover standby environment where the entire database can be restored and then selectively extract the application data needed to revert the changes in the Control-M data. *See [XML utilities](#) (such as *emdef* or *exportdeftable*) for extracting an Enterprise Manager specific set of application data from a given environment.*

Using Automation API you can export most of the Control-M objects either in XML or in JSON format. See the description of the Deploy service in the Automation API documentation. This procedure provides a scripted way to create a point-in-place backup of Control-M Folders, Calendars, Connection Profiles, Site Standards, etc, which allows you to restore previous versions of job definitions using standard application functionality rather than going back to an older date backup.

As another option for maintaining a backup for the purpose of having the ability to restore application data from some past date, it is a best practice to use Enterprise Manager utilities to export the application data such as job/calendar definitions, security, and then have this data available for restore in the primary environment. These utilities can be executed periodically without having to bring the Control-M application down. *See the description of the 'util' command and XML utilities in the Control-M Utilities Guide.*

Best Practice for Performing Backups with BMC Provided PG in Control-M HA Mode

When running backup utilities, or any administrative utilities, in a HA environment, you need to execute the utilities in the active environment or primary server and not on the secondary server that is not active. In a HA

configuration the server that is not active does not have access to the active process of the primary configuration and, therefore, is not able to connect to the services that it needs to complete the utilities.

Additional information can be found in knowledge base [000311604](#)

Best Practices for Disaster Recovery

In a Disaster Recovery (DR) scenario, the assumption is that the primary site, where Control-M components are installed, and a database server component that services Control-M components, will cease to exist and cannot be recovered. A failover Control-M environment with similar configuration is needed to take over all processing from the point of failure. Some loss of data may be tolerated when manual recovery of data takes a short time. DR configuration has a much larger tolerance in most of the business continuity requirements than HA configuration.

The Control-M architecture is a scalable, enterprise-wide solution with distributed components residing all across a connective network. The distributed Control-M components can serve as redundant entities to support failover or they can serve for load balancing in capacity planning. Only Control-M/Enterprise Manager supports distributed architecture for the majority of its add-ons. A distributed architecture of Control-M/Enterprise Manager can provide both scalability and redundancy. Building several instances of Control-M/Enterprise Manager facilitates a federated architecture for both GUI servers and for Web Servers. The Web Server of a Control-M/Enterprise Manager instance automatically load balances a user request to the least utilized GUI server. In a failover scenario, the Web Server recognizes when a given Control-M/Enterprise Manager instance becomes unavailable and routes requests only to active GUI servers. A third-party Load Balancing application can be utilized to load balance the traffic across several Web Servers. In Control-M 9.0.21, GUI servers are stateless and no “stick bit” is required on the Load Balancing application.

When analyzing DR needs, each Control-M component should be considered individually — whether it provides redundancy functionality or mainstream functionality. An entire Control-M/Enterprise Manager installation with its distributed components is a single functional unit, and when replicating for DR needs, the entire Control-M/Enterprise Manager installation should be replicated as a unit. Control-M in a single landscape can support only one Control-M/Enterprise Manager installation. Even if you chose to use the **One Install** (where all components are installed under the same user account), for DR purposes, the Enterprise Manager and Control-M/Server must be evaluated independently and treated as separate components.

The database server, which is an integral part of Control-M components, is a separate critical application. It is either internal to Control-M and integrated from the time of installation (as in the case of BMC PostgreSQL), or it is external to Control-M and installed in a different account or location to service the needs of the Control-M

application (as in the of existing Oracle, external PG, and MS SQL configurations). The management of external database applications is outside the scope of this document; we will focus only on the data that is application-specific and on the connection properties that allow Control-M to utilize the services of an external database server.

The discussion of database availability was covered above in the Database Solutions section. Regardless of the option that you select to achieve database recovery, the procedure you select must follow the guidelines of the specific database vendor; simply copying files from the database and restoring them in a new location breaks all data consistency and integrity aspects of a database server and is not supported. Control-M simplifies the basic backup/restore of the database operation by providing utilities that wrap the vendor-provided functionality to invoke backup/restore procedures for both internal installation and somewhat limited versions for external installation of database servers.

The primary function of the Control-M/Server component is to manage enterprise-wide workload for execution of scheduled and event-based tasks that most of the time are time sensitive and mission critical. The sequences of events that occur in the Control-M/Server are transactional, and therefore, recovery of the database must be at the point of failover to ensure that no job is executed twice after the recovery. Additionally, the integrity of the data is critical for the application to resume functioning from the time of recovery. As a result, our recommendation is to implement a standby database solution, such as Oracle DataGuard, Replication Servers, hot backups with immediate log shipping, and Control-M/Server HA solution to achieve recoverability with full data integrity that span distant network sites. Control-M/Server is a high transaction processing application and requires an infrastructure to support such demands. See *Best Practices for Full Control-M Databases Backups and Database Solutions on how to implement various database availability solutions*.

The rest of this section will focus on best practices for hot backup and restore as the sole solution for Disaster Recovery that does not involve any additional utilities or third-party components that are not packaged in the product. When using PostgreSQL, BMC recommends that Control-M hot backup procedures provide asynchronous replication, which is encapsulated in the product for Oracle, PostgreSQL in all releases, and Microsoft S SQL Server databases.

In order to accomplish a failover environment, the following must be achieved

1. Select a failover platform that will facilitate the same database vendor's ability to connect and restore from the hot backup performed on the primary environment.
2. Recover Control-M application on a failover environment.

3. Recover database connection properties for external database servers.
4. Recover and restore the database hot backup that was created on the primary site.
5. Reconfigure local data that resides in the database.

Selecting a failover platform is fully dependent on the ability to restore a hot backup created on the primary platform. The hot backup functionality is provided by the database vendor. Control-M does not implement the backup; rather it wraps the database vendor's own functionality for performing hot backup and restore. Therefore, the failover platform has to be a supported platform where Control-M and the database client are on the list of supported platforms. *For Product Compatibility and Availability, see the Control-M Customer Support website or Release Notes for your specific product and version.*

As an example, where the primary environment is RedHat Linux with an Oracle database, you can build a failover environment on an AIX® platform because the AIX platform is compatible for Control-M with an Oracle database. In all cases, the database vendor must be the same for primary and failover environments. In an environment where the database is PostgreSQL, the primary and failover platforms must be the same because the backup also includes the PostgreSQL binary environment. Therefore, in order to restore a PostgreSQL database in a hot backup environment, the failover must be the same platform as the primary.

Recovering a Control-M component in a failover environment can be accomplished only through a new Control-M installation performed from scratch in the failover environment. As specified above, file-by-file copy or a basic file backup and restore of the Control-M directory tree is not supported. A full installation of Control-M, including the database, should be performed in a failover environment, independently of any aspect of the primary environment. If the database is external, then it's up to the database services (DBA) to either provide a separate database within close proximity to the failover installation or to provide connection properties to a database server that will service the failover environment and will guarantee that the network latency will not impact the throughput of Control-M.

When installing the failover Control-M, ensure that the following:

- OS account owner for Control-M user is the same on both environments
- Path to Control-M account is the same from root directory on both environments

These parameters will ensure that jobs that need to run as Control-M owner will not require any modification on the failover environment. For PostgreSQL database configuration, the owner, owner ID, group, group ID, DBO, DBA, DBO passwords, and path must be identical.

In a scenario where the continuity of primary database services is guaranteed — and when Control-M on the failover environment will be connecting back to the primary database services — then the connection properties in a failover environment will need to be configured to point to the database services that guarantee database continuity. In this scenario, the following is needed to configure connection properties:

- The configuration includes changing the connection properties configuration files found in interface file `tnsnames.ora`, `interface`, and `postresq.conf` files.
- Modifying DBO user and password, database name, and instance name

Refer to documentation for the 'restore_host_config -reconf_db' utility to facilitate these changes for both Enterprise Manager and Control-M/Server

When the failover Control-M application has access to the database server, the data restore operation can proceed. Prior to running a hot restore, make sure that all the archived logs and the base backup files from the primary database location are safely copied to the failover machine or are accessible to the failover Control-M account prior to the restore. You will be prompted by the restore database command to specify the base backup file that was brought from the primary site as the source for database restore. You also will be prompted to specify the archived log directory that contains all the archived log files. The hot restore procedure will restore the base copy of the backup and then incorporate all the archived log files into the final configuration of the database. In an actual DR scenario where the last archived log file may not have been closed and shipped over to the failover site, there will be some discrepancy between the restored data and final state of some jobs in the primary Control-M environment. In this case, manual intervention may be required to reconcile the data between the two sites. For a DR scenario that is controlled, you can ensure that archived logs are closed and shipped prior to bringing down the primary site. In such a controlled environment, the restore operation will ensure that the data is restored to the point of failover.

Recovery to the point of failure is critical for Control-M/Server because a loss of a single transaction can make a difference between knowing if the last job was submitted to the agent for execution or it is still waiting to be sent. Therefore, recovery to point of failure on Control-M/Server is essential. However, recovery to point of failure on Enterprise Manager is much less critical because all the active data on the Enterprise Manager will automatically synchronize from its corresponding Control-M/Server or Control-M for z/OS®. The data that may be of possible concern is the definition data. As a result, the more convenient option for data recovery for Enterprise Manager may be to utilize the **util** export and import utility to extract all the data from the primary Enterprise Manager on a daily basis and then import it to the failover environment in case of disaster. As mentioned previously, the **util** export can be executed when Enterprise Manager is up and running and it does not require the overhead of handling log shipping and maintenance.

After the database is restored in a failover environment — and prior to starting the Control-M application — you need to ensure that the database contains local environment information and is adjusted to reflect the current failover environment. More specifically:

- The local hostname that may be hardcoded in Enterprise Manager in the System Parameters and in other configuration files is updated to reflect the new environment. *See the description of “restore_host_config” in the Control-M Utilities Guide for the steps needed to make the imported data reflect the new local environment.*
- Control-M/Server maintains full pathname to executables in the **CMS_SYSPRM** and **CMS_CMNPRM** tables. The **CMS_SYSPRM** and **CMS_CMNPRM** tables need to be updated to reflect the new paths for the **LOGDIR**, **PROCLIB**, and **EXE_PATH** fields. This step can be skipped if both primary and failover site have the same path from the root directory to the Control-M account. *See the description of the restore_host_config utility in the Control-M Utilities Guide.*

Step by step procedures for DR test and recover are provided in the following Knowledge Base articles:

[000211840](#) - Documents for Control-M/Server with Oracle / MSSQL

[000244906](#) – Documents for Control-M/Server with Postgres

For More Information

For more information about application continuity, please visit www.bmc.com/control-m.

Business runs on IT. IT runs on BMC Software.

Business runs better when IT runs at its best. That's why more than 20,000 IT organizations – from the Global 100 to the smallest businesses – in over 120 countries rely on BMC Software (NASDAQ: BMC) to manage their business services and applications across distributed, mainframe, virtual and cloud environments. With the leading Business Service Management platform, Cloud Management, and the industry's broadest choice of IT management solutions, BMC helps customers cut costs, reduce risk and achieve business objectives. For the four fiscal quarters ended June 30, 2012, BMC revenue was approximately \$2.2 billion.

Where to Get the Latest Product Information

To view the latest BMC Software documents, visit the BMC Customer Support page at http://www.bmc.com/support_home . BMC Software distributes printed copies of flashes, technical bulletins, and release notes with most product shipments, as indicated on your shipping list. In addition, all notices are available on the Customer Support page, including any notices that BMC Software issues after you receive your product shipment. You will not receive new notices by mail. However, by subscribing to proactive notification,

you can receive e-mail messages that direct you to those notices. For more information about proactive notification, refer to the Customer Support page.

BMC delivers software solutions that help IT transform digital enterprises for the ultimate competitive business advantage. We have worked with thousands of leading companies to create and deliver powerful IT management services. From mainframe to cloud to mobile, we pair high-speed digital innovation with robust IT industrialization—allowing our customers to provide amazing user experiences with optimized IT performance, cost, compliance, and productivity. We believe that technology is the heart of every business, and that IT drives business to the digital age.

BMC – Bring IT to Life