The TRELLIS™ Real-Time Infrastructure Optimization Platform

Technical Bulletin

Disaster Recovery

The information contained in this document is subject to change without notice and may not be suitable for all applications. While every precaution has been taken to ensure the accuracy and completeness of this document, Vertiv assumes no responsibility and disclaims all liability for damages resulting from use of this information or for any errors or omissions. Refer to other local practices or building codes as applicable for the correct methods, tools, and materials to be used in performing procedures not specifically described in this document.

Technical Support Site

If you encounter any installation or operational issues with your product, check the pertinent section of this manual to see if the issue can be resolved by following outlined procedures.
Visit https://www.VertivCo.com/en-us/support/ for additional assistance.

TABLE OF CONTENTS

1 OVERVIEW

In the event of a disaster, it is critical that your business IT and technology infrastructure systems are able to be recovered and continue running seamlessly. A disaster recovery plan creates company-specific policies and procedures to minimize the negative effects of the disaster and quickly resume mission-critical functions. Not having a disaster recovery plan in place not only puts your business at risk of incurring high financial costs after a disaster, but it also risks damage to your reputation, clients and customers. A major piece of mapping out your disaster recovery plan involves collecting information and capturing an accurate picture of your data centers company-wide. By utilizing the features of the *Trellis*™ Real-Time Infrastructure Optimization platform, you are provided a one-stop location to house your data center, device and power connection information to feed into your existing disaster recovery plan. The *Trellis*™ platform is customizable to fit your specific plan and provides configuration for a production site and a disaster recovery site. Full systems are set up in a data center at each site, with the production site system actively running and the disaster recovery site system turned off. If the production site fails during a disaster, all data stored in the *Trellis*™ platform is restored to the disaster recovery site, appliances are reconfigured to point to the disaster recovery system and platform applications are restarted. When a new production site is up and running, all of the information can then be transferred from the disaster recovery site back to the primary production location.

NOTE: Contact your Professional Services partner to customize the process to your specific environment.

When planning a disaster recovery strategy, reference your company's business continuity plan, which should indicate the key metrics of Recovery Point Objective (RPO) and Recovery Time Objective (RTO) for business processes. The RPO is the maximum acceptable amount of data loss measured in time, while the RTO is the maximum acceptable amount of time for restoring a computer system, network or application. Using the *Trellis* platform and the most recent backup data, the RTO is typically within four to six hours. This is dependent, however, on the deployment size, RPO and when the last backup occurred (typically within not more than 24 hours). While the production site is down, appliances continue to collect real-time data for up to seven days. When the disaster recovery site is active and appliances are reconfigured, the RPO for real-time data is 15 minutes.

The following diagram shows the typical disaster recovery plan set up using the *Trellis* platform.

Figure 3.1 Typical Disaster Recovery Solution with the *Trellis*™ Platform

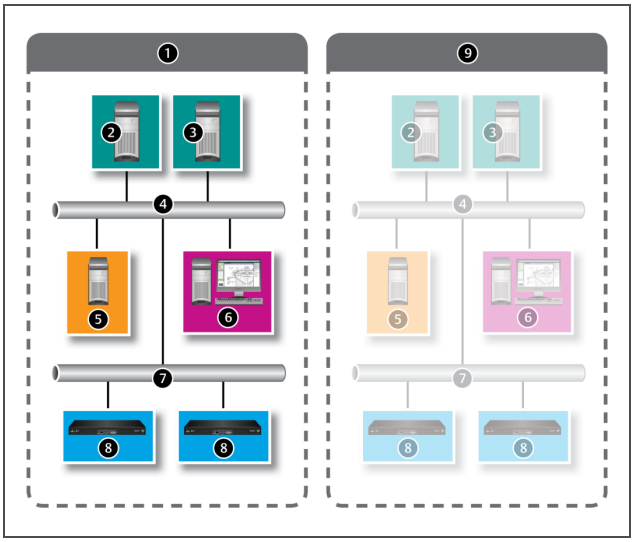

Table 3.1 Disaster Recovery Solution Descriptions

ITEM	DESCRIPTION
1	Production site data center hosting the active production system and the live instance of the *Trellis* platform. An active system is currently processing requests for the platform applications. In a normal scenario, the production site system is active and the disaster recovery site system is in cold standby mode.
2	Production site back machine. This is the back end of the software platform running at the production site.
3	Production site front machine. This is the front end of the software platform running at the production site.
4	The production site intranet and general purpose network.
5	Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) server for any authentication services/systems.
6	Client workstation.
7	Your infrastructure or management network.
8	The Avocent® Universal Management Gateway appliance and its embedded *Trellis*™ Intelligence Engine are monitored data center infrastructure devices that can be included at your production site and other data centers. The *Trellis*™ Intelligence Engine provided with the *Trellis*™ Site Manager module is also a monitored device.
9	Disaster recovery site data center hosting the backup of the production site system in cold standby mode (mirrors the production site). Cold standby mode indicates that the disaster recovery site is currently turned off. If a disaster happens, the disaster recovery scenario is when the production system is unavailable or moves to cold standby mode and the disaster recovery site is now the active system.
9.2	Disaster recovery site back machine in cold standby mode. This is the backup for the production site back machine.
9.3	Disaster recovery site front machine in cold standby mode. This is the backup for the production site front machine.
9.4	The disaster recovery site intranet and general purpose network.
9.5	Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) server for any authentication services/systems.
9.6	Backup Client workstation.
9.7	Your backup infrastructure or management network.
9.8	The Avocent® Universal Management Gateway appliance and its embedded *Trellis*™ Intelligence Engine are monitored data center infrastructure devices that can be included at your production site and other data centers. The *Trellis*™ Intelligence Engine provided with the *Trellis*™ Site Manager module is also a monitored device.

2 GETTING STARTED

Although it is assumed that your production and disaster recovery sites are already in place, a Professional Services partner will handle advanced configuration steps involved in your disaster recovery plan with the *Trellis*™ platform. However, there are additional steps that you need to perform to ensure your systems are ready.

NOTE: To ensure proper deployment of your disaster recovery plan, you must be an administrator or have experience with the platform software and hardware architecture. You should also be familiar with the processes and procedures in the *Trellis*™ Real-Time Infrastructure Optimization Platform Backup and Restore Procedure Technical Bulletins for Microsoft® Windows® and Red Hat® Enterprise Linux®.

2.1 Mapping Your Current System

Information about the current production site configuration must be captured to understand how to properly set up the disaster recovery site. Using the Appendices on page 17, record the current configurations for the Avocent® Universal Management Gateway appliances and the front and back machines at the primary production site. Also list and record the location of all artifacts and scripts located on the primary production system. As you begin building the disaster recovery site, configurations at that location can be listed there as well.

2.1.1 Prerequisites

The following prerequisites must be completed in order to configure your production and disaster recovery sites:

- Confirm the production system is regularly backed up or a snapshot of the system is stored at an offsite location.
- Create and maintain a change log that includes all changes at the production site and the supporting directory structure.
- Create a Network File Share (NFS) for both environments where Recovery Manager (RMAN) backups are saved or generated.
- Synchronize both the production and disaster recovery system servers to the same designated time server.
- Ensure any *Trellis*™ Intelligence Engines, either accompanying the Site Manager module or the Avocent® Universal Management Gateway appliance are registered with certificates in both systems and monitoring each data center infrastructure environment.

NOTE: All Avocent Universal Management Gateway appliances and the *Trellis*™ Intelligence Engines use the *Trellis*™ platform front machine Fully Qualified Domain Name (FQDN) instead of the IP address. When appliance registration is complete, the configuration finishes with a script that allows the appliance to use the FQDN to communicate to the production system. In the event of a disaster, this FQDN will be updated to point to the disaster recovery system instead.

- Ensure both systems have their own entitlements.
- Ensure the disaster recovery site system has the same application modules and device counts as the production site system.
- Ensure the production system is being actively monitored and the monitoring application alerts you in the event of system damage or failure.

2.2 Preparing the Production Site System

The following procedures will prepare your *Trellis* platform production site system for disaster recovery.

2.2.1 Using the RMAN tool

The database backup and recovery tool, RMAN, must be enabled before backups can be saved, generated or stored. RMAN is installed automatically on the database server in the /u01/trellis directory. The factory default setting for tool functionality is disabled. The RMAN tool must be enabled. After the RMAN tool has been enabled, you can configure and verify the RMAN backup files on your production site.

To enable the RMAN tool:

1. As **oracle**, log into the production site back machine.
2. For a Linux system, execute **/u01/trellis/configure.sh** to open the Configuration Menu.

   -or-

   For a Windows system, click *Start – Command Prompt*, then right-click and select *Run as Administrator*. At the prompt, enter **c:\u01\trellis\configure** to open the Configuration Menu.
3. Select option **1** for RMAN configuration.
4. Select option **1** and press **Enter** to enable the RMAN backup.
5. Accept the default Fast Recovery Area (FRA) location and press **Enter** or enter a new location.
6. Enter **YES** to enable the RMAN backup. This executes the first database backup and schedules the backup job. Full backups of the database include: the complete contents of all data files of the database, the control file, archived redo log files and the server parameter file. With these files, you can perform a complete recovery.

NOTE: The default backups run daily at midnight.

To configure and verify the generated RMAN backup files:

1. As **root** on the production site back machine, create an NFS share between the production and disaster recovery systems (for example, the /u05/backup directory location). The database RMAN backups are generated and saved in this directory, which must have read and write access by both the oracle user and oinstall group.
2. As **oracle**, log into the production site back machine.
3. For a Red Hat® Enterprise Linux®, execute **/u01/trellis/configure.sh** and select option **1** to open the Configuration Menu.

   -or-

   For a Windows system, click *Start – Command Prompt,* then right-click and select *Run as Administrator*. At the prompt, enter **c:\u01\trellis\configure** and select option **1** to open the Configuration Menu.
4. Select option **1** (Modify Retention Policy) and press **Enter** to accept the retention policy.
5. On the Configuration Menu, select option **4** (Reports), then select option **2** (Backup History) to view the status of each backup.
6. When the backup is complete, navigate to */u05/backup* and verify the ORCL folder was created during the backup.

NOTE: The RMAN backup folder on the NFS drive and its children must have permissions set to oracle:oinstall in Red Hat® Enterprise Linux®.

2.2.2 Staging artifacts

When mapping out your production site, you listed and recorded any artifacts and scripts being used on your system (see the Appendices on page 17 for more information). These artifacts and scripts also need to be staged and prepared for transfer over to the disaster recovery system. Ensure that you have completed the following steps:

- Create staging for all deployed artifacts on the server file system before deployment. Ensure that all artifacts and scripts are deployed from the same location on the production system.
- Create a subdirectory for each deployment using a timestamp to uniquely identify the change set and keep the disaster recovery site in sync. Custom artifacts should be placed in the /u05/customdeployment folder.
- Create a folder for each set of artifacts you wish to deploy using the following format for the file system deployment folder: YYYY-MM-DD-HH-MM.

While some artifacts are automatically pushed to and restored at the disaster recovery site, some need to be manually deployed and maintained. The following table lists artifacts, where they will be deployed at the disaster recovery site and notes if the deployment is automatic or manual.

Table 3.2 Artifact Deployment

ARTIFACT TYPE	DEPLOYED LOCATION	DEPLOYMENT PROCESS
Stored procedures	Database	Automatically restored to disaster recovery system during database restoration.
Business processes or composites	Database	Automatically restored to disaster recovery system during database restoration.
Java Enterprise archives and web archives	File system	Must be manually deployed to the disaster recovery system and maintained.
Java message queues	File system and database	Must be manually deployed to the disaster recovery system and maintained. As a best practice, when defining a storage location, always choose the database over the file system to ensure automatic failover of business transactions that are in process.
Service bus configuration	File system	Must be manually deployed to the disaster recovery system and maintained.

2.3 Preparing the Disaster Recovery Site System

NOTE: The following steps are performed by Professional Services.

To configure the disaster recovery site:

CAUTION: It is assumed that the disaster recovery system is not created on the same physical network as the production system. If both systems are on the same physical network, but the production system is not shut down, the systems may destabilize and cause permanent damage.

1. On the production site back machine, log in as **oracle** and generate the TrellisConfiguration.zip file. For Red Hat® Enterprise Linux®, execute **/u01/trellis/configure.sh** and select option **2**.

   -or-

   For Windows, execute **c:/u01/configure.cmd** and select option **2**.

2. Export the *Trellis*™ platform configuration.
3. Save the *TrellisConfiguration.zip* file to a specified location.

4. Copy the *TrellisConfiguration.zip* file generated in step 3 from the production site back machine to the /home/oracle/ folder in the disaster recovery back machine.
5. Run the *Trellis* platform software installation on the disaster recovery back machine.
6. After installation in the disaster recovery back machine is complete, copy the TrellisConfiguration.zip file from the disaster recovery back machine to the disaster recovery front machine in the /home/oracle folder.
7. Run the *Trellis* platform software installation on the disaster recovery front machine.
8. Log into the *Trellis* platform to verify the platform software is fully installed and running on the disaster recovery system.
9. On the production site back machine, navigate to */u01/app/oracle/product/12.1.0.2/dbs* and copy the spfileorcl.ora file to the same directory on the disaster recovery back machine.

NOTE: Whenever a change is made to the *Trellis* platform, the spfileorcl.ora file must be copied over from the production back machine to the disaster recovery back machine.

10. Configure customer-specific software settings.
11. Register and configure the *Trellis*™ Intelligence Engines, either accompanying the *Trellis*™ Site Manager module or the Avocent® Universal Management Gateway appliance, to monitor the data center infrastructure.
12. Configure all entitlements and certificates on the disaster recovery system.
    a. For entitlements, ensure the production system and disaster recovery system have the same application modules and device counts.
    b. For certificates, ensure that all Avocent® Universal Management Gateway appliance, Intelligence Engine and third party integration certificates are copied from the production system to the disaster recovery system.
13. Place the disaster recovery system in cold standby mode.
    a. As **oracle** on the front machine, enter **/etc/init.d/trellis stop** to stop the platform.
    b. Repeat step a on the back machine.
    c. As **root** on the back machine, enter **/etc/init.d/oracle stop** to stop the Oracle database.
    d. Turn off the *Trellis*™ disaster recovery machines.

NOTE: When the production system is active, the disaster recovery system must be in cold standby mode.

3 HANDLING A DISASTER

If a disaster occurs and the production site is destroyed or damaged, it is critical to bring the disaster recovery site online as quickly as possible. As soon as your production site monitoring application alerts you of damage or system failure, contact your Professional Services partner.

3.1 Bringing the Disaster Recovery Site System Online

NOTE: Professional Services will complete the following procedures.

3.1.1 Restoring the databases

As soon as the state of the production site is confirmed, the disaster recovery system must be started and the databases must be recovered, reconfigured as needed and restarted.

To activate the disaster recovery site system:

1. If the production site is not a complete loss and some equipment is still accessible, ensure all production site machines are in cold standby mode.
    a. As **oracle** on the production system front machine, enter **/etc/init.d/trellis stop** to stop the platform.
    b. Repeat step a on the production back machine.
    c. As **root** on the production system back machine, enter **/etc/init.d/oracle stop** to stop the Oracle database.
    d. Turn off the *Trellis*™ production site machines.
2. Turn on the disaster recovery site front and back machines.
3. If the disaster recovery site machines are in a different time zone than the production site machines, change the disaster recovery site machines to match the time zone of the production site machines.

To restore the customer database from a backup:

NOTE: It is assumed that the *Trellis*™ platform and Oracle databases were not yet started on the disaster recovery site machines. If the databases are running, enter /etc/init.d/trellis stop on the front and back machines to stop the platform, then log in as root on the back machine and enter /etc/init.d/oracle stop to stop the Oracle database.

1. For Red Hat® Enterprise Linux®, as **root**, log onto the disaster recovery site back machine.

    -or-

    For Windows, as an administrator log onto the disaster recovery site back machine.

2. Keeping the directory structured as it is already for Red Hat® Enterprise Linux®, enter **rm -f /u02/app/oracle/oradata/orcl/*** to remove the files from the /u02/app/oracle/oradata/orcl directory.

    -or-

    For Windows, browse to the *c:\u01\app\oracle\oradata\orcl* directory and remove the files on this directory.

3. For Red Hat® Enterprise Linux®, enter **rm -f /u02/app/oracle/oradata/orcl/orcl/*** to remove the files from the other orcl directory.

    -or-

    For Windows, browse to the *c:\u01\app\oracle\oradata\orcl\orcl* directory and remove the files on this directory.

4. For Red Hat® Enterprise Linux®, still as the **root** user on the disaster recovery site back machine, enter **/etc/init.d/oracle start** to start the Oracle database.

   -or-

   For Windows, start the OracleServiceORCL service.

5. As **oracle**, log into the disaster recovery site back machine.
6. For a Red Hat® Enterprise Linux®, execute **/u01/trellis/configure.sh** and select option **1** to open the Configuration Menu.

   -or-

   For a Windows system, click *Start – Command Prompt*, then right-click and select *Run as Administrator*. At the prompt, enter **c:\u01\trellis\configure** and select option **1** to open the RMAN Configuration Menu.

7. Wait for the menu items to be updated and select option **1** (Restore Database) to display the list of available database backups to restore.
8. Enter the date and time for the restore. You can copy the date and time (for example, 27/07/2016 03:30:42) from one of the available database backups to restore, paste the text after the Restore to date prompt and press **Enter**. At the restore database confirmation prompt, enter **YES** and press **Enter**.
9. Open another session (also as **oracle** on the back machine), enter **tail –f400 /u02/app/oracle/diag/rdbms/orcl/orcl/trace/alert_orcl.log** to see the restore in progress. The Configuration Menu appears when the restore operation is complete. At this point, you can exit the Configuration Menu.

NOTE: After the restore is complete, you can adjust time zones to your desired time, if needed. If the disaster recovery site machines are in a different time zone than the production site machines, change the disaster recovery site machines to match the time zone of the production site machines

10. As **oracle** on the disaster recovery site back machine, enter **cd /u01/trellis/support** to access the support directory on the back machine.
11. For Linux, remove the old hostname from the OID tables by entering ./removeOIDoldhostname.sh on the disaster recovery site back machine.

    -or-

    For Windows, remove the old hostname from the OID tables by entering removeOIDhostname.cmd on the disaster recovery site back machine.

12. Review the removeOIDoldhostname.log generated in the /u03/logs directory and ensure references to the production site host servers are removed.
13. As **oracle**, enter **sqlplus / as sysdba @searchandreplace.sql** to replace old references to the front server hostname or IP address.
    a. In the value for search_string, enter the hostname of the front production machine and press **Enter** (hsv-tr-12180.systemtest.com, for example).
    b. In the value for replacement_string, enter the value of the hostname of the front disaster machine and press **Enter** (hsv-tr-08134.systemtest.com, for example).
    c. Review the replace.sql file created in the current directory.
14. As **oracle**, enter **sqlplus / as sysdba @replace.sql** to execute the replace script, then enter **exit** to quit SQL.

To run the disaster recovery site system after backups are restored:

1. If the Oracle database is not already running, for Red Hat® Enterprise Linux®, enter **/etc/init.d/oracle start** as the **root** user on the disaster recovery site back machine to start the Oracle database.

   -or-

   For Windows, start the OracleServiceORCL service.

2. As **oracle** on the disaster recovery site back machine, enter **/etc/init.d/trellis start** to start the platform database.
3. Repeat step 2 on the disaster recovery site front machine.
4. Perform a quick check on your systems to ensure they are running properly.

To verify *Trellis*™ platform functionality:

1. Log into the *Trellis* platform software, click the avatar drop-down arrow, select About and verify you are running the correct version.
2. From the menu bar, select the Portfolio pivot bar icon and verify current data is available.

3.1.2 Reconfiguring the appliances and TRELLIS™ Intelligence Engines

Trellis™ Intelligence Engines, either accompanying the *Trellis*™ Site Manager module or the Avocent® Universal Management Gateway appliance, use an FQDN to communicate with production site systems. After a disaster, the FQDN must be updated so that the appliances communicate with the disaster recovery system instead.

NOTE: During the disaster recovery process, it is not uncommon for the appliances to be in a non-responding state.

To bring the appliances running Intelligence Engine version 4.6.0.21 (or higher) back online:

1. Log into the *Trellis* platform software, select the Monitoring pivot bar icon, select *Data Collection Engine* and validate all data collection engines are still registered.
2. Using SSH or PuTTY, log into the Avocent® Universal Management Gateway appliance using the username **admin** and password **password**, and then enter **/mss/engine/bin** to navigate to that directory.

   -or-

   Using SSH or PuTTY, log into the Red Hat® Enterprise Linux® or Ubuntu operating system where the *Trellis*™ Intelligence Engine is installed and navigate to the /usr/bin directory.

3. For the appliance, enter **./mss-run ConfigUpdate.sh** to run the Reset Target (reset_Trellis_target) utility.

   -or-

   For the *Trellis*™ Intelligence Engine running in Red Hat® Enterprise Linux® and the appliance, enter **./mss_run ConfigUpdate.sh** to update the *Trellis*™ Address utility.

4. Enter option **4** to update the *Trellis* platform address.
5. Enter the IP address of the platform, FQDN or DNS alias of the disaster recovery front machine, then enter **No** to leave the port default unchanged. Upon completion of the utility, all services are up and running.
6. For the appliance, browse to the */mss/engine/conf/MssEngine.ini* and the */mss/mssengine/config/mssengine_config.ini* files to verify the new IP address or FQDN.

   -or-

For the *Trellis*™ Intelligence Engine running in Red Hat® Enterprise Linux® or Ubuntu, browse to the */etc/mss/conf/MssEngine.ini* and */etc/mss/mssengine/config/mssengine_config.ini* files to verify the new IP address or FQDN.

7. Repeat steps 1-6 for every Avocent Universal Management Gateway appliance or *Trellis*™ Intelligence Engine in your system.
8. Log back into the *Trellis* platform software, select the Monitoring pivot bar icon and click *Data Collection Engine*.
9. Validate that all data collection engines are responding and real-time data collection is working properly.

To bring appliances back online:

NOTE: Avocent® Universal Management Gateway appliances running firmware versions earlier than 2.9.0.25 do not have the Reset Target utility pre-installed. Contact Technical Support for assistance to acquire the appropriate version of the Reset Target utility.

1. Log into the appliance using a file transfer utility (such as WinSCP) and navigate to the var/home directory.
2. Copy the ResetTrellisTarget.tar.gz file over to the /var/home directory on the appliance.
3. Log out of the copy file utility.
4. Using SSH or PuTTY, log into the Avocent Universal Management Gateway appliance and navigate to the /var/home directory.
5. Enter cd /var/home and tar –zxvf ResetTrellisTarget.tar.gz to unpack the tarball, create the ResetTrellisTarget directory and unpack the files.
6. Open the ResetTrellisTarget directory and run the utility by entering the following commands:

   cd ResetTrellisTarget

   chmod 755 reset_Trellis_target

   ./reset_Trellis_target <IP or FQDN of the disaster recovery site front machine>

7. When the utility finishes executing, check the */mss/engine/conf/MssEngine.ini* and the */mss/mssengine/mssengine_config.ini* files to ensure they have the new IP address or FQDN.
8. Repeat steps 1-7 for every Avocent® Universal Management Gateway appliance in your system.
9. Log back into the *Trellis*™ platform software, select the Monitoring pivot bar icon and click *Data Collection Engine*.
10. Validate that all data collection engines are responding and real-time data collection is working properly.

NOTE: After the system front and back machines, databases and appliances are running, restored and reconfigured, the disaster recovery site should be functional and normal backups are now initiated on this system. Your Professional Services partner will notify you that the disaster recovery system is now available.

3.2 Running Final System Verifications

After your Professional Services partner has notified you that the disaster recovery system is online, run the following verification tests on your system to ensure it is working properly.

To verify your system:

1. Log into the *Trellis* platform software.
2. Click the Monitoring pivot bar icon, select *Data Collection Engine* and validate all data collection engines are registered and responding.

3. Click the Administration pivot bar icon, select *User Management* and *Roles* and then validate the Authentication and Roles settings are configured as expected.
4. Click the Process pivot bar icon and validate the Business Process Manager settings are configured as expected.
5. Click the Process pivot bar icon and verify you can access Business Process Manager workflows.
6. Click the Reports pivot bar icon and check reporting functionality.
7. When these verifications are complete, notify your users they can access the platform with full functionality.

4 RE-ESTABLISHING THE PRODUCTION SITE

At this point in the disaster recovery process, the original production site is no longer available and the disaster recovery site has taken its place as a temporary production site until a new site is built to replace the one that was lost. In this scenario, the disaster recovery site now contains the active systems running the *Trellis*™ platform. When a new production site is ready, everything stored on the disaster recovery site must be transferred to that location.

Because the steps to restore the database from the disaster recovery site to the production site are similar, the following steps to restore the database from the original production site, the same comments apply.

4.1 Bringing the New Production Site System Online

NOTE: Professional Services will complete the following procedures.

4.1.1 Restoring the databases

When the new production site is ready, the disaster recovery system must be shut down and the databases must be restored at the new production site.

To activate the new production site system:

1. Verify that the backup on the disaster recovery system is up-to-date and the Flash Recovery Area (FRA) directory is accessible. If any changes have been made since the last backup, a manual backup is required.
2. Ensure all disaster recovery site machines are in cold standby mode.
    a. As **oracle** on the disaster recovery system front machine, enter **/etc/init.d/trellis stop** to stop the platform.
    b. Repeat step a on the back machine.
    c. As **root** on the disaster recovery system back machine, enter **/etc/init.d/oracle stop** to stop the Oracle database.
    d. Turn off the *Trellis* platform disaster recovery site machines.
3. Turn on the new production site front and back machines.
4. If the new production site machines are in a different time zone than the disaster recovery machines, change the new production site machines to match the time zone of the disaster recovery site machines.

To restore the customer database from the backup on the disaster recovery system:

NOTE: It is assumed that the *Trellis* platform and Oracle databases were not yet started on the new production site machines. If the databases are running, enter /etc/init.d/trellis stop on the front and back machines to stop the platform, then log in as root on the back machine and enter /etc/init.d/oracle stop to stop the Oracle database.

1. As **root**, log onto the new production site back machine.
2. Keeping the directory structured as it is already, for Linux, enter **rm -f /u02/app/oracle/oradata/orcl/*** to remove the files from the /u02/app/oracle/oradata/orcl directory.

    -or-

    For Windows, browse to the c:\u01\app\oracle\oradata\orcl directory and remove the files from this directory.
3. For Red Hat® Enterprise Linux®, enter **rm -f /u02/app/oracle/oradata/orcl/orcl/*** to remove the files from the other orcl directory.

    -or-

    For Windows, browse to the *c:\u01\app\oracle\oradata\orcl\orcl* directory and remove the files on this directory.

4. For Red Hat® Enterprise Linux®, still as the **root** user on the disaster recovery site back machine, enter **/etc/init.d/oracle start** to start the Oracle database.

   -or-

   For Windows, start the OracleServiceORCL service.

5. As **oracle**, log into the new production site back machine.
6. For a Red Hat® Enterprise Linux®, execute **/u01/trellis/configure.sh** and select option **1** to open the RMAN Configuration Menu.

   -or-

   For a Windows system, click *Start – Command Prompt*, then right-click and select *Run as Administrator*. At the prompt, enter **c:\u01\trellis\configure** and select option **1** to open the Configuration Menu.

7. Wait for the menu items to be updated and select option **1** (Restore Database) to display the list of available database backups to restore.
8. Enter the date and time for the restore. You can copy the date and time (for example, 27/07/2016 03:30:42) from one of the available database backups to restore, paste the text after the Restore to date prompt and press **Enter**. At the restore database confirmation prompt, enter **YES** and press **Enter**.
9. Open another session (also as **oracle** on the back machine), enter **tail –f400 /u02/app/oracle/diag/rdbms/orcl/orcl/trace/alert_orcl.log** to see the restore in progress. The Configuration Menu appears when the restore operation is complete. At this point, you can exit the Configuration Menu.

NOTE: After the restore is complete, you can adjust time zones to your desired time, if needed.

10. As **oracle** on the new production site back machine, enter **cd /u01/trellis/support** to access the support directory on the back machine.
11. For Red Hat® Enterprise Linux®, remove the old hostname from the OID tables by entering **./removeOIDoldhostname.sh** on the new production site back machine.

    -or-

    For Windows, remove the old hostname from the OID tables by entering **removeOIDhostname.cmd** on the new production site back machine.

12. Review the removeOIDoldhostname.log generated in the /u03/logs directory and ensure references to the disaster recovery site host servers are removed.
13. As **oracle**, enter **sqlplus / as sysdba @searchandreplace.sql** to replace old references to the front server hostname or IP address.
    a. In the value for search_string, enter the hostname (hsv-tr-12180.systemtest.com, for example) of the front production machine and press **Enter**.
    b. In the value for replacement_string, enter the value of the hostname (hsv-tr-08134.systemtest.com, for example) of the front disaster machine and press **Enter**.
    c. Review the replace.sql file created in the current directory.
14. As **oracle**, enter **sqlplus / as sysdba @replace.sql** to execute the replace script, then enter **exit** to quit SQL.

To run the new production site system after backups are restored:

1. If the Oracle database is not already running, for Red Hat® Enterprise Linux®, enter **/etc/init.d/oracle start** as the **root** user on the disaster recovery site back machine to start the Oracle database.

   -or-

For Windows, start the OracleServiceORCL service.

2. As **oracle** on the new production site back machine, enter **/etc/init.d/trellis start** to start the platform database.
3. Repeat step 2 on the new production site front machine.
4. Perform a quick check on your systems to ensure they are running properly.

To verify *Trellis*™ platform functionality:

1. Log into the *Trellis*™ platform software, click the avatar drop-down arrow, select *About* and verify you are running the correct version.
2. From the menu bar, select the Portfolio pivot bar icon and verify current data is available.

4.1.2 Reconfiguring the data collection engines

After the new production site is up and running, the FQDN must be updated so that the data collection engines communicate with the new site system instead of the disaster recovery site system.

NOTE: As information is transferred from the disaster recovery site to the new production site, it is not uncommon for the data collection engines to be in a non-responding state.

To bring the appliances running Intelligence Engine version 2.9.0.25 (or higher) back online:

1. Log into the *Trellis* platform software, click the Monitoring pivot bar icon, select *Data Collection Engine* and validate all data collection engines are still registered.
2. Using SSH or PuTTY, log into the Avocent® Universal Management Gateway appliance using the username **admin** and password **password**, and then enter **/mss/engine/bin** to navigate to that directory.

   -or-

   Using SSH or PuTTY, log into the Red Hat® Enterprise Linux® or Ubuntu operating systems where the *Trellis*™ Intelligence Engine is installed and navigate to the /usr/bin directory.

3. For the appliance, enter **./mss-run ConfigUpdate.sh** to run the Reset Target (reset_Trellis_target) utility.

   -or-

   For the *Trellis*™ Intelligence Engine running in Red Hat® Enterprise Linux® and the appliance, enter **./mss_run ConfigUpdate.sh** to update the *Trellis*™ Address utility.

4. Enter option **4** to update the *Trellis* platform address.
5. Enter the IP address of the platform, FQDN or DNS alias of the new production site front machine, then enter **No** to leave the port default unchanged. Upon completion of the utility, all services are up and running.
6. Browse to the *mss/engine/conf/MssEngine.ini* and */mss/mssengine/config/mssengine_config.ini* files to verify the new IP address or FQDN.

   -or-

   For the *Trellis*™ Intelligence Engine running in Red Hat® Enterprise Linux® or Ubuntu, browse to the */etc/mss/config/mssengine.ini* and */etc/mss/config/mssengine_config.ini* files to verify the new IP address or FQDN.

7. Log back into the *Trellis* platform software, click the Monitoring pivot bar icon and select *Data Collection Engine*.
8. Validate all data collection engines are responding and real-time data collection is working properly.

To bring the appliances running Intelligence Engine versions prior to 2.9.0.25 back online:

NOTE: Avocent® Universal Management Gateway appliances running firmware versions earlier than 2.9.0.25 do not have the Reset Target utility pre-installed. Contact Technical Support for assistance acquiring the appropriate version of the Reset Target utility.

1. Log into the appliance using a file transfer utility (such as WinSCP) and navigate to the var/home directory.
2. Copy the ResetTrellisTarget.tar.gz file over to the /var/home directory on the appliance.
3. Log out of the copy file utility.
4. Using SSH or PuTTY, log into the Avocent Universal Management Gateway appliance and navigate to the */var/home* directory.
5. Enter cd /var/home and tar –zxvf ResetTrellisTarget.tar.gz to unpack the tarball, create the ResetTrellisTarget directory and unpack the files.
6. Open the ResetTrellisTarget directory and run the utility by entering the following commands:

   cd ResetTrellisTarget

   chmod 755 reset_Trellis_target

   ./reset_Trellis_target <IP or FQDN of the new production site front machine>

7. When the utility finishes executing, check the /mss/engine/conf/MssEngine.ini and the /mss/mssengine/mssengine_config.ini files to ensure they have the new IP address or FQDN.
8. Repeat steps 1-7 for every Avocent® Universal Management Gateway appliance in your system.
9. Log back into the *Trellis*™ platform software, click the Monitoring pivot bar icon and select *Data Collection Engine*.
10. Validate that all data collection engines are responding and real-time data collection is working properly.

NOTE: After the system front and back machines, databases and appliances are running, transferred and reconfigured, the new production site should be functional and normal backups are now initiated on this system. Your Professional Services partner will notify you that the new production site system is now available.

4.2 Running Final System Verifications

After your Professional Services partner has notified you that the new production system is online, run the following verification tests on your system to ensure it is working properly.

To verify your system:

1. Log into the *Trellis*™ platform software.
2. Click the Monitoring pivot bar icon, select *Data Collection Engine* and validate all data collection engines are registered and responding.
3. Click the Administration pivot bar icon, select *User Management* and *Roles* and then validate the Authentication and Roles settings are configured as expected.
4. Click the Process icon and validate that Business Process Manager settings are configured as expected.
5. Click the Process icon and verify you can access Business Process Manager workflows.
6. Click the Reports icon and check reporting functionality.
7. When these verifications are complete, notify your users they can access the platform with full functionality.

This page intentionally left blank

5 APPENDICES

Appendix A: Site Configuration Records

Information about your original production site configuration must be captured to understand how to properly set up the disaster recovery site. Using the following tables, record the current configurations for the Avocent® Universal Management Gateway appliances, the *Trellis*™ Intelligence Engines running under Red Hat® Enterprise Linux® or Ubuntu and the front and back machines at the original production site. Also list and record the location of all artifacts and scripts on the production system. As you begin building the disaster recovery site, configurations at that location can also be listed.

NOTE: The first row of each table is provided as an example of configuration information and space is provided after the table for you to note any additional information.

Table 4.1 Original Production Site Configuration Information

SYSTEM	GLOBAL NAME	IP ADDRESS	SUBNET MASK	DEFAULT GATEWAY
Front machine running the *Trellis*™ platform	TrellisFrontMachine-XYZ.acme.com	10.10.10.11	255.255.255.0	10.10.10.1

Additional Configuration Notes:

Table 4.2 Certificates, Scripts and Artifacts on the Production System

ITEM NAME	PRODUCTION SYSTEM LOCATION	NOTES
Stored procedure	YYYY-MM-DD-HH:MM	Staged for deployment

Additional Configuration Notes:

Table 4.3 Appliance Configuration Information

NAME	NETWORK INTERFACE 0	SUBNET MASK	DEFAULT GATEWAY	NETWORK INTERFACE 1	SUBNET MASK	DEFAULT GATEWAY
DC1-UMG-001.acme.com	10.10.10.101	255.255.255.0	10.10.10.1	10.20.10.101	255.255.255.0	10.20.10.1

Additional Appliance Configuration Notes:

Table 4.4 Disaster Recovery Site Configuration Information

SYSTEM	GLOBAL NAME	IP ADDRESS	SUBNET MASK	DEFAULT GATEWAY
Front machine running the *Trellis*™ platform	TrellisFrontMachine-XYZ.acme.com	10.10.10.11	255.255.255.0	10.10.10.1

Additional Disaster Recovery Site Configuration Notes:

Table 4.5 Certificates, Scripts and Artifacts on the Disaster Recovery System

ITEM NAME	PRODUCTION SYSTEM LOCATION	DR SITE DEPLOYMENT LOCATION	NOTES
Stored procedure	YYYY-MM-DD-HH:MM	XYZ Database	Deployed to DR system

Additional Configuration Notes:

This page intentionally left blank


VERTIV.

VertivCo.com | Vertiv Headquarters, 1050 Dearborn Drive, Columbus, OH, 43085, USA

590-1542-501C