Hewlett Packard Enterprise

Application Performance Management

Version 9.40, Released August 2017

APM Redundancy - Best Practices

Published August 2017

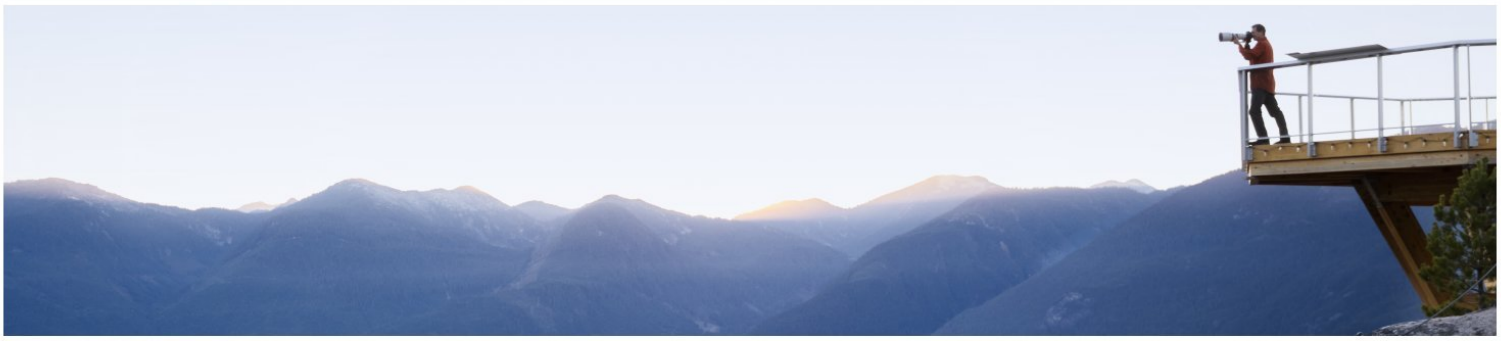

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

AMD and the AMD Arrow symbol are trademarks of Advanced Micro Devices, Inc.

Google™ and Google Maps™ are trademarks of Google Inc.

Intel®, Itanium®, Pentium®, and Intel® Xeon® are trademarks of Intel Corporation in the U.S. and other countries.

iPod is a trademark of Apple Computer, Inc.

Java is a registered trademark of Oracle and/or its affiliates.

Microsoft®, Windows®, Windows NT®, Windows® XP, and Windows Vista® are U.S. registered trademarks of Microsoft Corporation.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/group/softwaresupport/search-result?keyword=.

This site requires an HPE Passport account. If you do not have one, click the **Create an account** button on the HPE Passport Sign in page.

Support

Visit the HPE Software Support website at: **https://softwaresupport.hpe.com**

This website provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software Support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HPE Passport user and sign in. Many also require a support contract. To register for an HPE Passport ID, go to **https://softwaresupport.hpe.com** and click **Register**.

To find more information about access levels, go to:
https://softwaresupport.hpe.com/web/softwaresupport/access-levels

HPE Software Integrations and Solutions

Visit the Integrations and Solutions Catalog at **https://softwaresupport.hpe.com/km/KM01702731** to explore how the products in the HPE Software catalog work together, exchange information, and solve business needs.

Contents

Chapter 1: Introduction

HPE Application Performance Management (APM) is a suite of software that acts as a performance dashboard to present a comprehensive view of the network, software and system operations of a corporate data center. This software includes software designed to keep your business healthy by monitoring applications across traditional, mobile, virtual and cloud environments. APM provides insight into every transaction for quick resolution of application issues, and helps reduce costs by giving you a common tool for pre-production and production. Application Performance Management improves application performance by monitoring end-user experience, and aligning IT performance with business goals. Detailed diagnostics and real-time topology-based analytics improve application quality

In this document we will look at four alternatives for APM redundancy: High Availability, Disaster Recovery, Distributed High Availability, and Business Continuity. Each of these four use cases will be covered in detail, including architectural alternatives, advantages and disadvantages of each possibility, and best practices.

There are two metrics we look at with each use case:

- **Recovery Point Objective (RPO)**, an industry term used to describe the amount of changed data a business is willing to lose in an outage.
- **Recovery Time Objective (RTO)**, an industry term used to describe the desired maximum down time.

Ideally, we are looking for a system that will provide us with zero RPO and zero RTO. This can be achieved using the Business Continuity use case, however the process for doing this requires a considerable amount of customization and is often not a viable alternative. The goal of this document is to help you choose the best solution for your data center.

Note: You should also consider redundancy for APM Data Collectors (BPM, RUM, Diagnostics, and SiteScope), as well as for UCDMB. These topics are outside the scope of this document.

Chapter 1: APM High Availability

Introduction

HPE APM can be set up in a High Availability configuration. Setting up High Availability is documented in the APM 9.30 Installation Guide – Appendix F: High Availability for APM (referred to in this document as "Appendix F").

Setting up High Availability involves load balancing a set of APM Gateway Servers and setting up a Failover APM Data Processing Server. Load balancing the APM Gateway Servers ensure that there is no single point of failure. If one server in not available there will still be one or more servers available to process incoming data and serve the users. Enabling High Availability on the Data Processing Server will cause the High Availability Controller to perform automatic failover if it detects compromised Data Processing Server services. If this event occurs, services will be assigned to the backup Data Processing Server.

Advantages

- APM High Availability is a fully-supported out-of-the-box solution provided by HPE.
- Of the redundancy options, APM High Availability is the simplest, least expensive solution.
- Recovery Point Objective should be zero. The system is not accessible while the primary DPS fails over to the backup DPS and the backup DPS's High Availability services are starting. However, metrics, and topology coming in from data collectors should be persisted and recoverable once the backup DPS is operational.
- Recovery Time Objective is lower than when using APM Disaster Recovery.

Disadvantages

- Recovery Time Objective is greater than zero. You can expect an RTO of 15-60 minutes, depending on your environment. This depends on how long it takes to start up the backup DPS processes.
- APM High Availability must be implemented in the same data center. This scenario does not take a complete data center outage into consideration. Therefore it should not be considered a true Disaster Recovery solution. APM Distributed High Availability can be implemented in two data centers, but this is also not to be considered a true Disaster Recovery solution due to the physical proximity requirements between the two data centers.
- APM High Availability acts as a single APM instance rather than two independent APM instances.

Architecture

High Availability for APM can be represented by the following diagram. In this implementation all servers are contained in the same primary data center. Up to three Gateway Servers are balanced by a local load

balancer. A primary and failover Data Processing Server are deployed in the same data center.

There is only one logical database server utilized in APM High Availability, so there is no need to synchronize metrics, and topology. Keep in mind this logical database may have its own High Availability system with multiple physical servers. Oracle RAC is a good example of this.

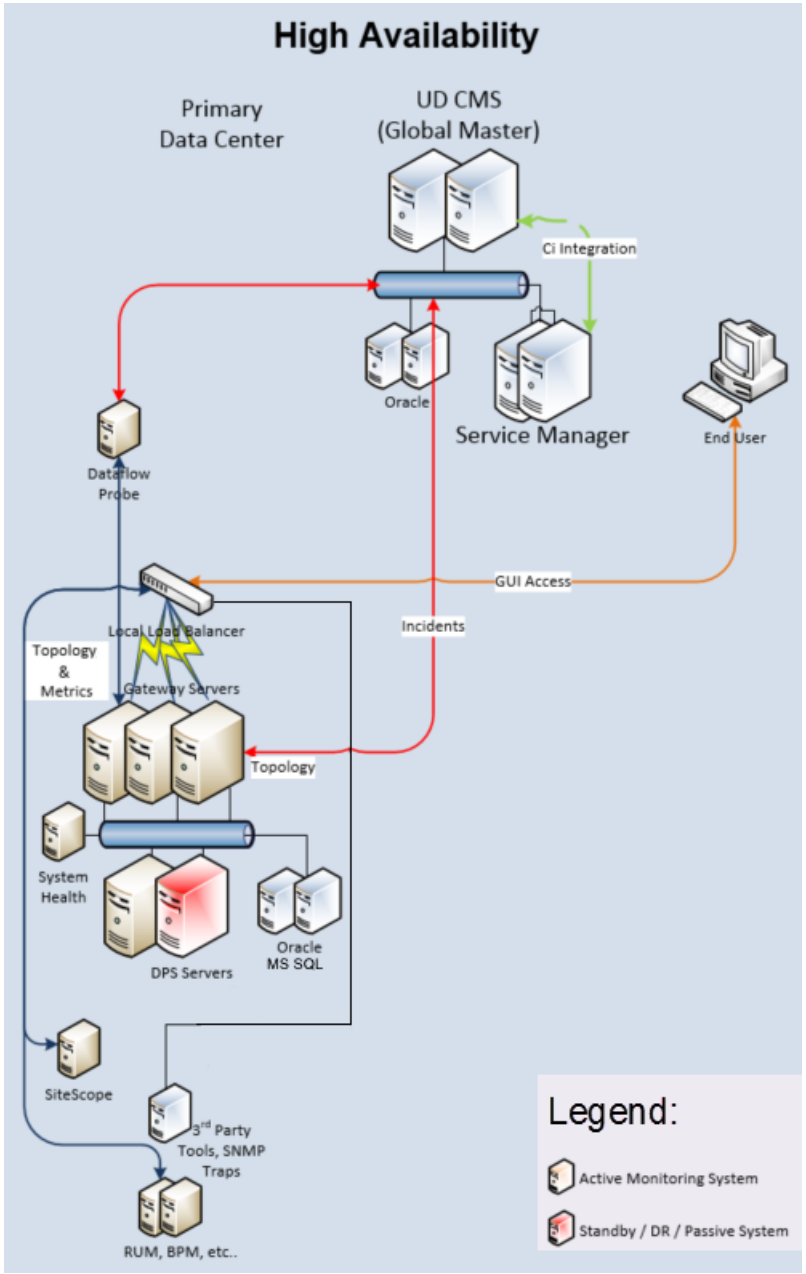

Best Practices

- Refer to the APM High Availability Fine Tuning Best Practices document.
- Automatic Data Processing Server Failover is not enabled out-of-the-box. To enable automatic detection

and fail-over of the APM Data Processing server's services, follow the instructions in the Configuring Automatic Failover section of Appendix F.

- The Primary Data Processing Server and the Failover Data Processing Server need to be comparable in terms of hardware, memory, network, and storage performance.
- Gateway servers need to be comparable in terms of hardware, memory, network, and storage performance.
- The first Data Processing Server that is started in APM deployment will become the primary DPS. The second DPS that is started can be assigned to act as a backup DPS.
- DPS services can be manually re-assigned using the JMX Console (see Appendix F), but an easier way to accomplish this is by using APM System Health. APM System Health has a user interface that allows for quick re-assignment of services.
- Automatic Data Processing Server Failover can be enabled by using the process in Appendix F, but an easier way to accomplish this is by using APM System Health. APM System Health has a user interface that will enable automatic failover:

Chapter 2: APM Distributed High Availability

Introduction

HPE APMcan be set up in a Distributed High Availability configuration. This is similar to the High Availability for APM setup, documented in the APM 9.30 Installation Guide – Appendix F: High Availability for APM (referred to in this document as “Appendix F”).

The key differences are that the Distributed High Availability setup requires the following:

- Deployment to two separate (distributed) data centers.
- A high speed network (<5ms latency) between the two data centers.
- A distributed APM database (see "Disadvantages" on the next page).

Setting up Distributed High Availability involves load balancing two sets of APM Gateway Servers (up to three servers in each set), and setting up a Primary and Failover APM Data Processing Server. Each set of Gateway Servers is deployed to one of two different data centers. The Primary and Failover Data Processing Servers are divided, with the Primary deployed in the first data center and the Failover deployed to the second data center. Load balancing the APM Gateway Servers in each data center ensures that there is no single point of failure. If one server in not available there will still be one or more servers available to process incoming data and serve the users.

Fail-over is a manual process; the links need to be reconfigured manually on the Global Load Balancer, and the DPS server needs to be manually failed over to the backup DPS – which is not active until each of the High Availability processes have been started.

There is only one logical database server utilized in APM High Availability, so no synchronization of metrics, and topology is needed. Keep in mind this logical database will have its own high availability system with multiple physical servers. Oracle RAC and SQL Server 2012 Always-On Cluster are examples of this.

Advantages

- APM Distributed High Availability is a fully-supported out-of-the-box solution provided by HPE.
- APM Distributed High Availability is a reasonably inexpensive solution to maintain.
- APM Distributed High Availability can be implemented in two different data centers. Because there are strict requirements on the data center distance and network latency, this cannot be considered a Disaster Recovery solution. This can be a good compromise for customers that have two data centers located close to each other (<400 km / 250 miles) with a dedicated high speed fiber connection between them.
- Recovery Point Objective should be close to zero. The system is not accessible while the primary DPS fails over to the backup DPS and the backup DPS’s High Availability services are starting. However, metrics and topology coming in from data collectors should be persisted and recoverable once the backup DPS is operational. A manual process is needed for fail-over, which may impact your recovery point

depending on how long it takes to execute this process.

- Recovery Time Objective is lower than when using APM Disaster Recovery.

Disadvantages

- Recovery Time Objective is greater than zero. Realistically we can expect an RTO of 30-60 minutes depending on your environment. This depends on how long it takes to manually fail over the Global Load Balancer and start up the backup DPS server's high availability processes.
- Fail over is a manual process. The links need to be reconfigured manually on the Global Load Balancer, and the DPS server needs to be manually failed over to the backup DPS – which is not active until each of the High Availability process have been started.
- APM Distributed High Availability acts as a single APM instance rather than two independent APM instances.
- Requires a distributed APM database (for example Oracle RAC Extended Distance Clusters or a Multi-Site SQL Server 2012 Always-On Cluster).
- Requires a very fast network connection (<5 millisecond network latency) between data centers. The two data centers must be within 400 km (250 miles) from each other, with dedicated fiber for the APM application between each data center (see Best Practices section below).

Architecture

The following diagram illustrates the concept of APM Distributed High Availability – this splits the APM High Availability solution between two data centers that have a very fast network link (<5ms latency) between them.

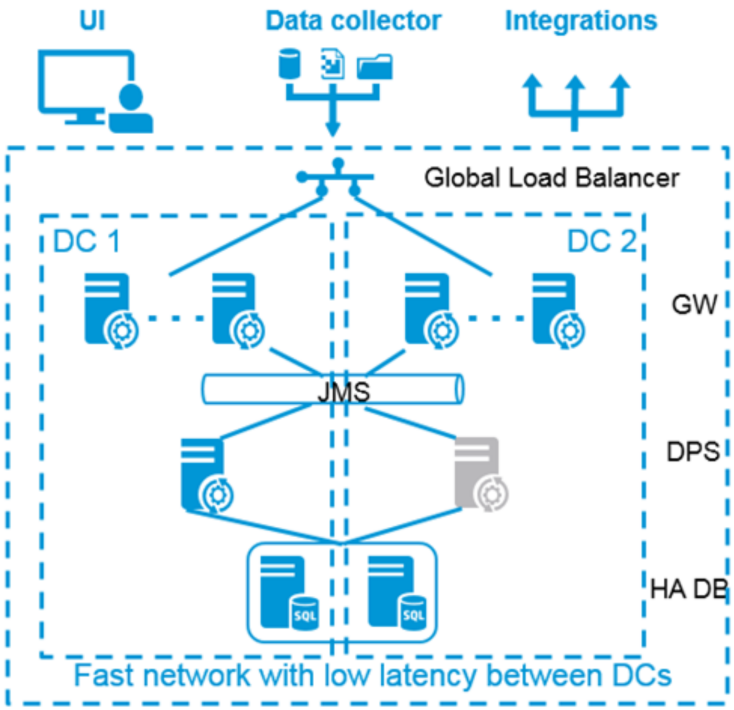

A more detailed implementation diagram is shown below. This shows the APM Gateway servers and APM Data Processing Servers split across the two different data centers. A database cluster is also housed with nodes in each data center. The APM Database acts as a single database clustered across two different physical locations.

The failover DPS is not active until failover occurs. Failover occurs manually by:

- Reconfiguring the Global Load Balancer to send traffic to the failover site.
- Manually failing over APM using either:
  - The JMX Console (instructions in Appendix F: Reassigning Services with JMX Console)
  - BAPM SM System Health

Up to three different Gateway Servers can be configured at each location.

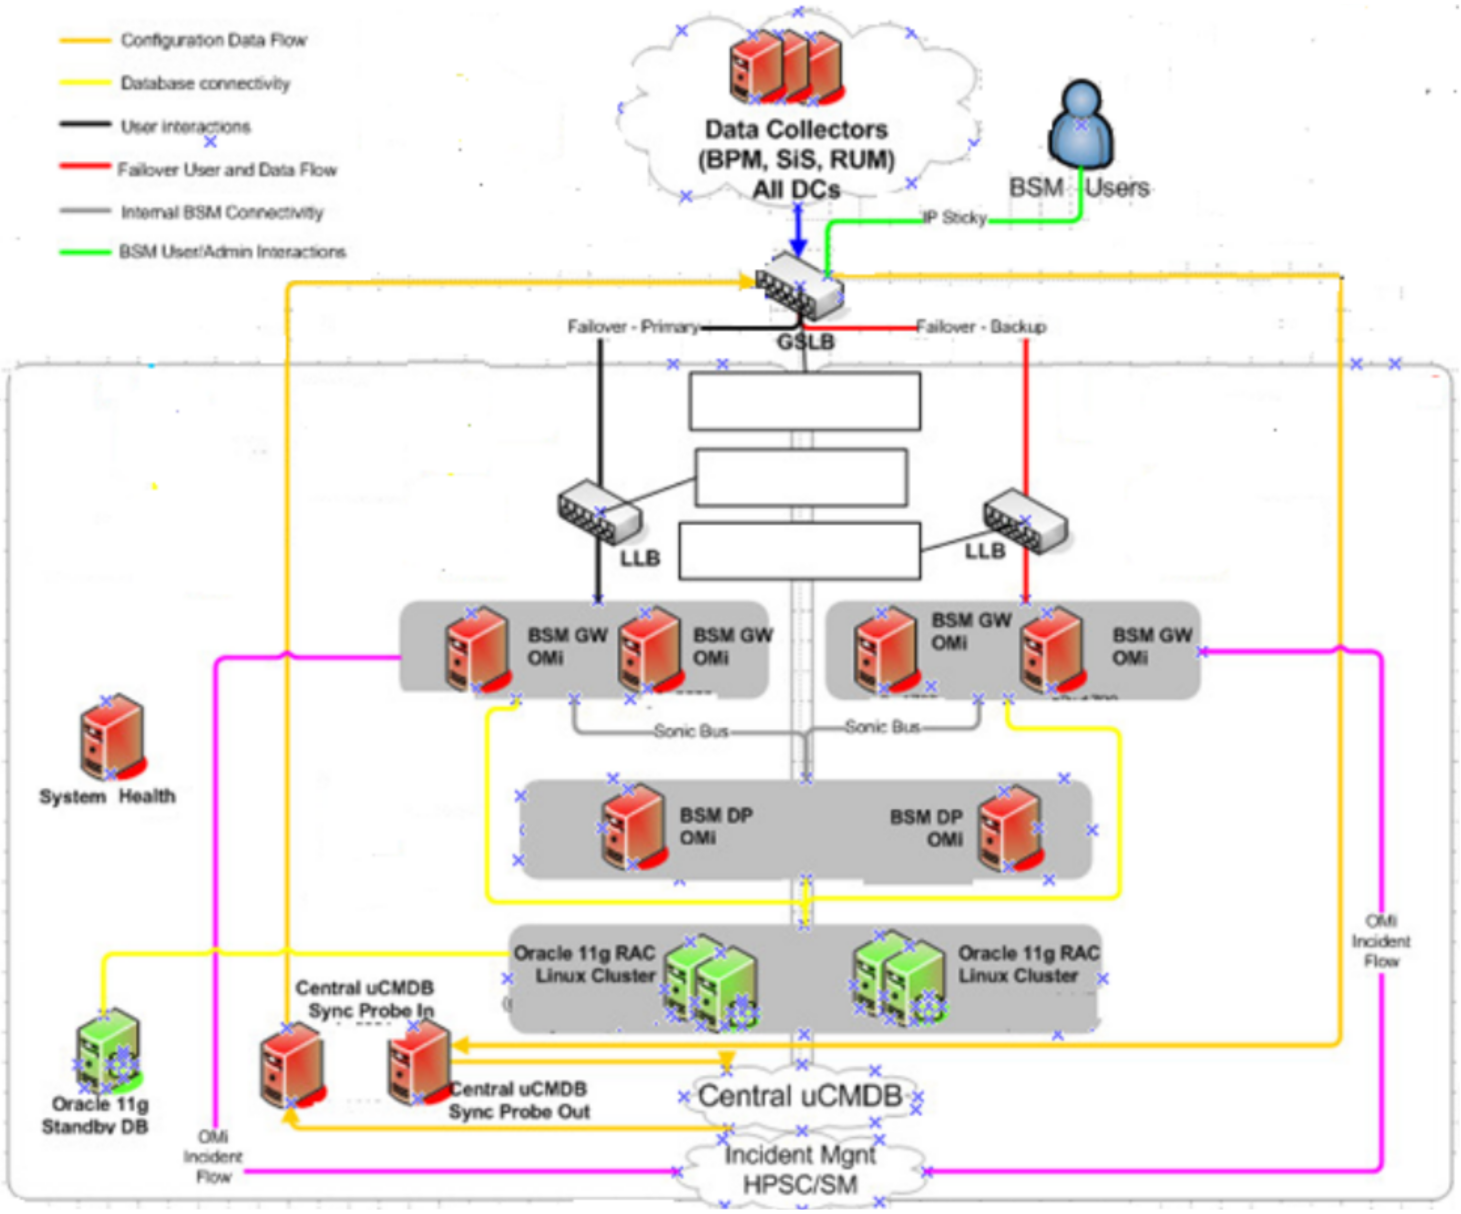

Best Practices

Refer to the APM High Availability Fine Tuning Best Practices document.

Best Practices for this solution are identical to the High Availability "Best Practices" on page 7. There are two additional best practices:

- Ensure you have a consistent, <5 millisecond connection between the two data centers. Note the following:
  - When connection speeds cannot exceed 5 milliseconds, based on the speed of light over fiber and certain guard bands for network delays, a maximum distance of 400 km (250 miles) between data centers should be assumed (see http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DCI/4-0/EMC/dciEmc/EMC_2.html).
  - This would assume dedicated fiber capable of carrying the communication bandwidth for the APM infrastructure between the two data centers.
  - It is important to benchmark the APM application and APM database network traffic between the two data centers during normal operations and during failover operations to determine how much fiber needs to be dedicated to the APM infrastructure.
- Document and practice the manual failover process. If failover is practiced regularly it becomes a routine operation rather than an operation that is unfamiliar during a crisis, which will reduce your Recovery Time Objective.

Chapter 3: APM Disaster Recovery

Introduction

Duplicate HPE APM instances can be set up in a Disaster Recovery configuration. Setting up Disaster Recovery is documented in the APM 9.30 Installation Guide – Appendix E: Disaster Recovery for APM (referred to in this document as "Appendix E").

Setting up Disaster Recovery involves setting up two complete APM systems: the **APM Production Instance** – with one to three Gateway Servers and one or two Data Processing Servers, and the **APM Failover Instance** – identical to the APM Production Instance (a second set of Gateway and Data Processing servers). In this configuration there are separate logical databases for the APM Production instance and the APM Failover Instance. The two logical databases in this solution are replicated using the database vendor's replication solution (for example Oracle Data Guard).

There are important notes in Appendix E that must be considered:

- Disaster Recovery involves manual steps in moving various configuration files and updates to the APM database schemas. This procedure requires at least one APM administrator and one database administrator who is familiar with the APM databases and schemas.
- There are a number of different possible deployment and configurations for APM. To validate that the disaster recovery scenario works in a particular environment, it should be thoroughly tested and documented. Contact HPE Professional Services to ensure best practices are used in the design and failover workflow for any disaster recovery scenario.
- A disaster recovery machine must use the same operating system and root directory as the original environment.

The manual steps involved can take from one to three hours, depending on:

- The complexity of your environment
- How well-rehearsed you are at performing the disaster recovery activity
- Whether the process in Appendix E has been specifically documented for your environment.

This document will cover the basic steps in Appendix E, and will provide:

- Additional guidance on alternative architecture, including the use of an external Configuration Management System (CMS)
- Guidance on how to create an environment that is conducive for trouble-free disaster recovery by:
  - Creating a Cleanup procedure that is customized to your environment
  - Practicing the process on a regular basis

Advantages

- APM Disaster Recovery is less expensive to maintain than APM Business Continuity, and still allows for recovery in a true disaster situation.

- APM Disaster Recovery is not restricted to a single data center.
- There is no restriction on the proximity between data centers or the speed of the communication link between the data centers.
- Recovery Time Objective can be reduced by following best practices.

Disadvantages

- APM Disaster Recovery is more expensive to maintain than either APM High Availability or APM Distributed High Availability.
- Recovery Point Objective is highest of the four alternatives (dependent on replication and speed of running clean-up procedure). Realistically we can expect an RPO of 2 – 24 hours depending on your environment.
- Recovery Time Objective is highest of alternatives (dependent on replication and speed of running clean-up procedure). Realistically we can expect an RTO of 2 – 24 hours depending on your environment.

Architecture

The architecture illustrated in Appendix E does not account for an external Configuration Management System (CMS) – or in this example a Universal Configuration Management Database (UCMDB). The following is a copy of the diagram from Appendix E:

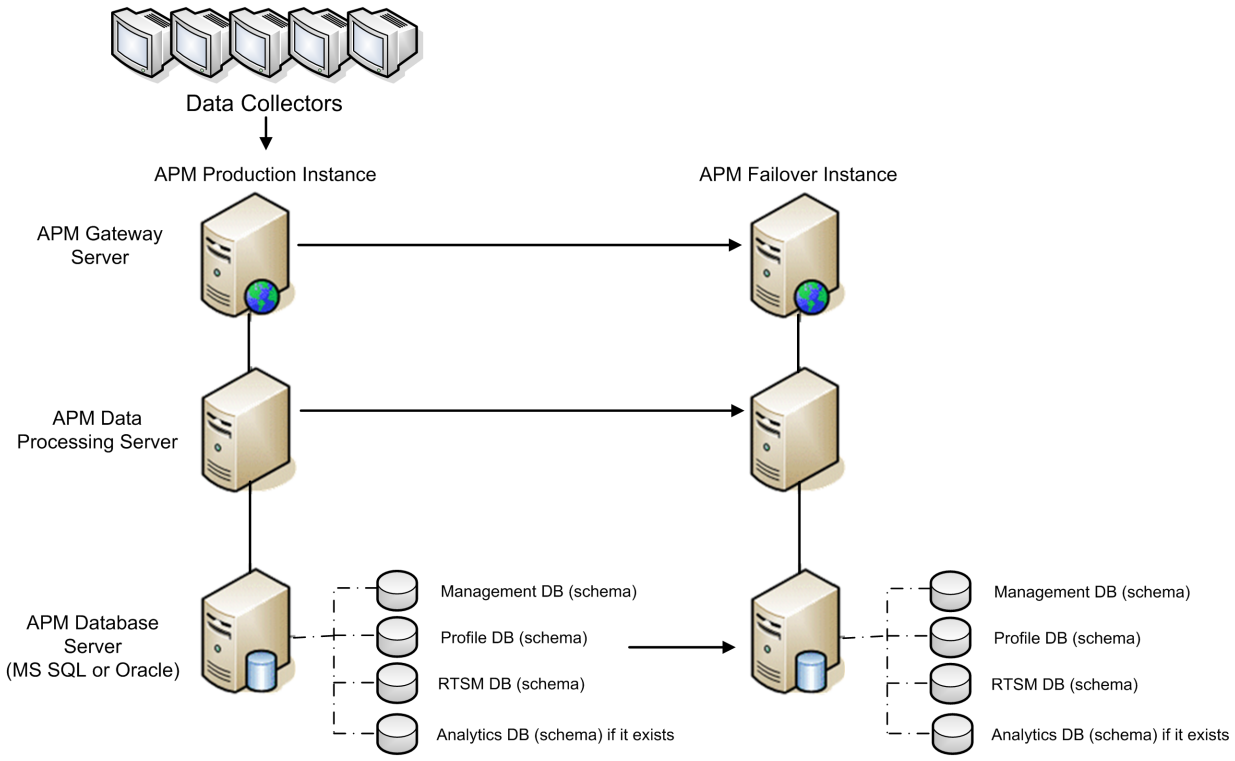

In the example above (with no external UCMDB) database replication is used to replicate all of the APM databases. In the next example, there is an external UCMDB as the topology system of record. Therefore, only the Management and Profile databases are replicated. Topology synchronization from the external UCMDB keeps the RTSM databases synchronized.

This example can be modified to accommodate for UCMDB synchronization. In this example there is data populating from the Universal CMDB's CMS database to each of the APM RTSM instances (both Production and Failover). This populates necessary CIs in the RTSM that are being discovered by Universal Discovery.

The Production APM Instance populates CIs back to the UCMDB database. This populates the CIs that are discovered by the APM Data Collectors.

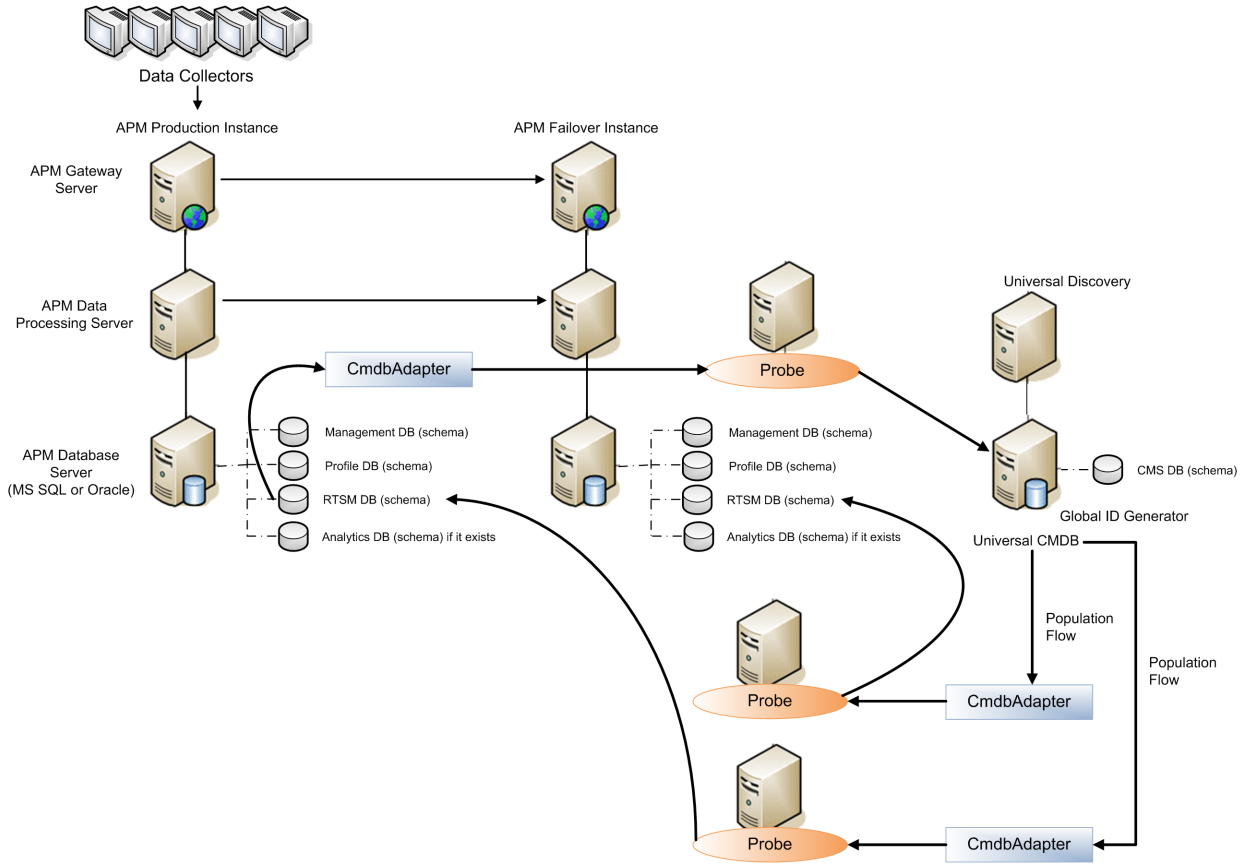

The environment in your data center may vary considerably from this, and accommodations may be needed for:

- Multiple staged UCMDB installations
- Data filtering of data populating the RTSMs from the UCMDB
- Data filtering of data populating the UCMDB from the Production RTSM
- The type of integration that is being used between each RTSM and the UCMDB (Population, Data Push, or Federation)

Create a Clean-up Procedure For Your Environment

Note that the failover procedure is a highly manual process; Appendix E describes the clean-up procedure in a five page section, which needs to be tailored for your environment.

The best ways to reduce the overall time it takes to fail over to the Disaster Recovery environment are to do the following:

1. **Create Your Own Clean-up Procedure.** Copy the clean-up procedure that is documented in Appendix E and make the appropriate changes needed for your environment. There are a number of parameters that need to have real values substituted, including:
   - <context value>
   - <new value>
   - <key>
   - NewDatabasehostname
   - NEWDatabaseServerName
   - NEWSID
   - OLDSID
   - NEW_UID_name
   - OLD_UID_name
   - NEW_port_name
   - OLD_port_name

   You should know these values, and create a document with these values populated. If you prepare this before you need to perform Disaster Recovery, you will avoid confusion and mistakes. You will also save time instead of needing to look up these values.
2. **Practice.** If you practice Disaster Recovery regularly, it becomes a routine operation rather than an operation that is unfamiliar during a stressful time. Disaster recovery can also be used to minimize downtime and risk during patch upgrades.

Best Practices

APM Disaster Recovery is not a trivial or completely automated process. The complexity of the system configuration and the high volume of data that is being synchronized presents challenges for even the most experienced administrators.

We recommend you do the following:

- Thoroughly review Appendix E.
- Modify the diagrams in Appendix E to conform with your environment.
- Modify the clean-up procedure in Appendix E to conform with your environment.
- Regularly practice disaster recovery during scheduled outages, and continue to modify your internal documentation to correct for changes in your data center.

Using these practices will create an environment where you can confidently recover from unforeseen problems.

Chapter 4: APM Business Continuity

Introduction

Duplicate APM instances can be set up in a Business Continuity configuration. Setting up Business Continuity is similar to setting up APM Disaster Recovery (documented in the APM 9.30 Installation Guide – Appendix E: Disaster Recovery for APM).

The key differences are:

- The two APM Database systems are separate and independent of each other.
- No data is replicated between the two APM Databases.
- All configuration, including data collector integrations, alert configurations, downtime configurations, saved user reports, and RTSM Packages must be manually entered on both systems.

Setting up Business Continuity involves setting up two complete APM systems:

1. **APM Production Instance**, with one to three Gateway Servers and one or two Data Processing Servers,
2. **APM Failover Instance**, identical to the APM Production Instance, with a second set of Gateway and Data Processing Servers.

In this configuration there are separate logical databases for the APM Production Instance and the APM Failover Instance; the two logical databases in this solution are not replicated.

Advantages

- APM Business Continuity is not restricted to a single data center.
- Recovery Point Objective of close to zero can be achieved.
- Recovery Time Objective of close to zero can be achieved.
- APM Business Continuity can be used to reduce the user load on Gateway Servers, as users can be set up to use the closest APM system.
- APM Business Continuity can be used to reduce the data entry load on Gateway Servers and Data Collectors, which may result in the ability to increase polling durations. For example, a single APM instance with 1 SiteScope running all monitors at 5 minute intervals can be split up across two APM instances, with two SiteScope servers running all monitors at 10 minute intervals. This effectively provides 5 minute polling intervals with only half the load on each APM instance.

Disadvantages

- APM Business Continuity is the most expensive solution to maintain by far, as each data collector, monitor, and configuration (alerts, saved User Reports, key performance indicators, health indicators, service level agreements) needs to be duplicated on each system

- APM Business Continuity requires discipline in managing all of the APM configurations (alerts, saved user reports, key performance indicators, health indicators, service level agreements) on two separate APM instances.
- APM Business Continuity requires duplication of most APM data collectors (SiteScope, Business Process Monitor, and Real User Monitor).
- Metrics from data collectors may not match each other precisely on both instances, since these are not synchronized to the millisecond. For example, Business Process Monitor #1 connected to APM#1 may report that a transaction failed. Several milliseconds later Business Process Monitor #2 connected to APM#2 may report that the transaction succeeded. This can cause concerns for systems that are closely audited.
- Because there is duplication of monitoring, each target system could have twice the amount of monitoring load on it. For example, instead of one data collector monitoring a target system once every 5 minutes, there will be two data collectors monitoring a target system once every 5 minutes. This can be accommodated for by increasing intervals – because there are two data collectors monitoring the target system, decision makers may agree to set the interval for each data collector to once every 10 minutes.
- APM Business Continuity will double the number of licenses needed for APM and each of its data collectors.
- APM Business Continuity is currently a customized integration between Service Manager, Operations Orchestration, and HPE APM instances. This requires the development of Groovy scripting to ensure that duplicate incidents are not opened in Service Manager, and that duplicate Runbooks are not generated in Operations Orchestration.
- APM Business Continuity is an approach that has been implemented in the field with limited success. The key challenge with this solution is the discipline it takes to manually ensure each instance is synchronized with the other. Duplicating configuration entries on two different systems is a process that requires a high amount of discipline and maturity for the team that manages the solution.
- The APM Business Continuity solution is not a fully tested, fully documented or fully supported approach.

Architecture

The architecture of APM Business Continuity is illustrated below:

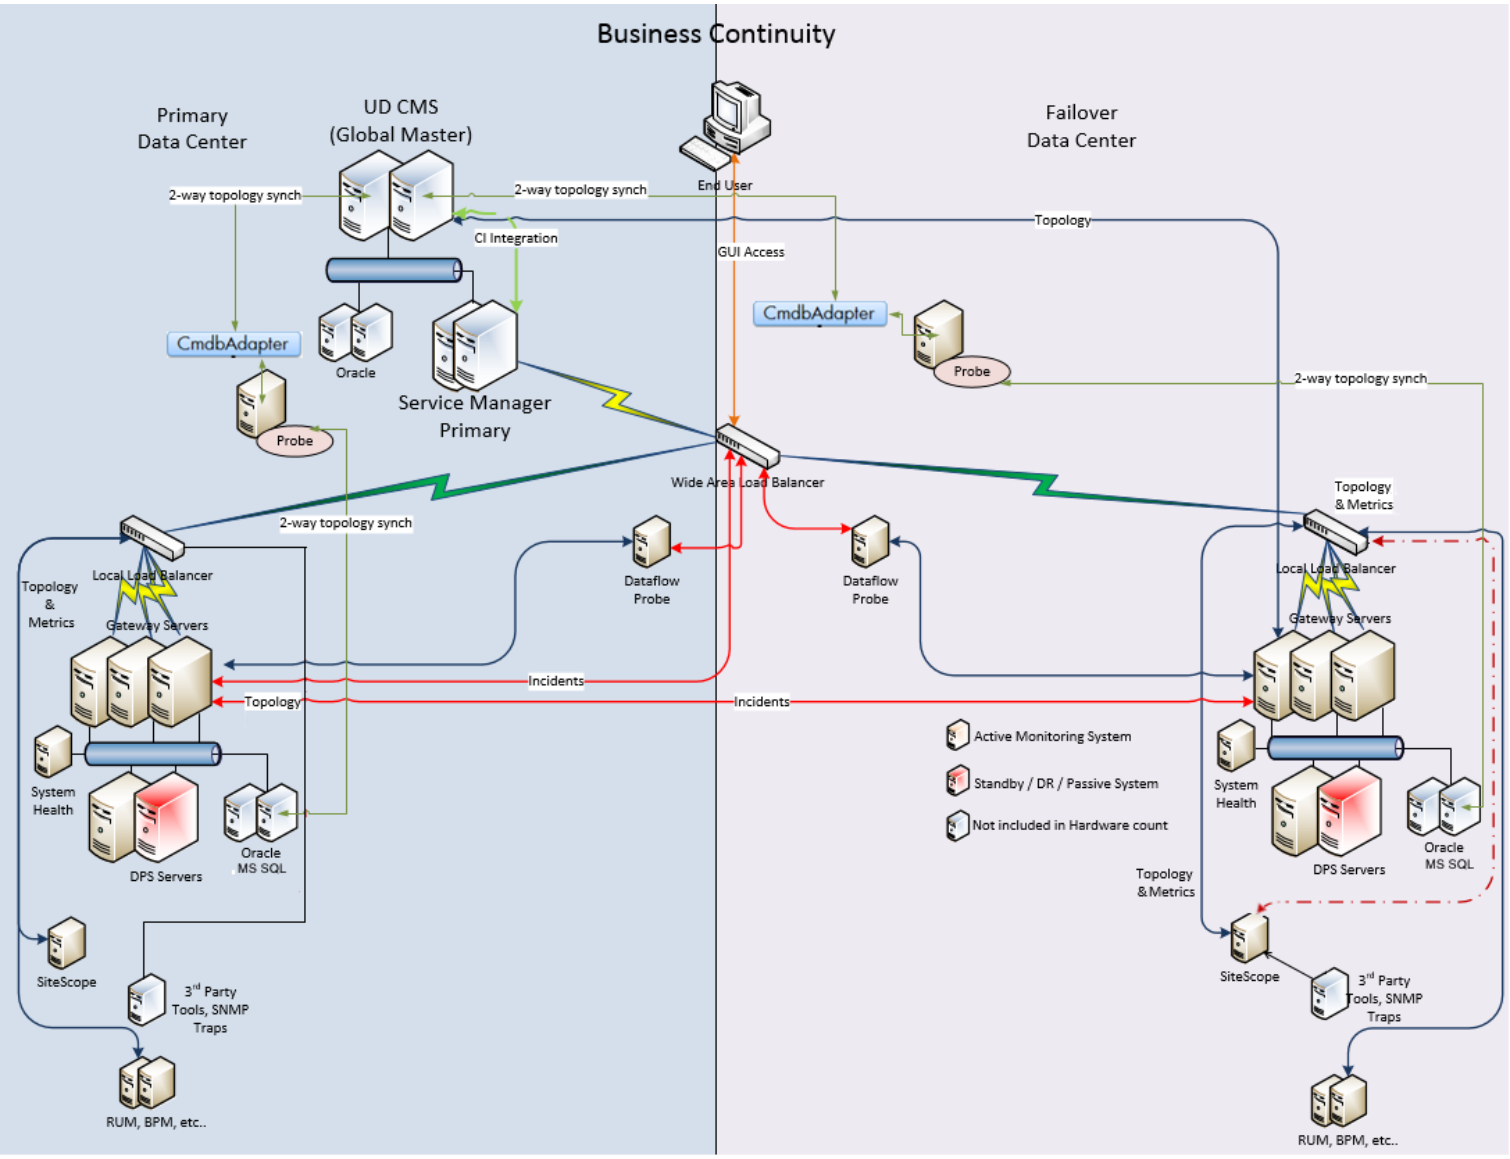

This architecture requires the following:

1. Two stand-alone implementations of APM (one in each Data Center).
2. Integration with UCMDB (Master for Global ID's).
3. Integration of both APM instances with Service Manager via wide area load balancer. This is currently a customized solution, where the first APM instance acquires the event, marks the event as the owner, and forwards it to the second APM . The second APM instance sees the same event, but only increments the counter and forwards the same event back to the first APM instance. Custom Groovy scripting is developed, and a CMA is assigned to each event to prevent duplication of incidents in Service Manager or duplication of Operations Orchestration flows being fired off.
4. One wide area load balancer to point users at the Data Center Tools in use.
   a. This requires a Virtual Internet Protocol (VIP) for each user interface.
   b. Service Manager VIP address would be the target for integrations.
   c. UCMDB VIP address would be the target for integrations.
   d. APM user VIP address would be the target (set to sticky for the current active APM local user VIP).

5. Two local area load balancers to point users and Data Collectors at the local APMs.
    a. This will require a VIP for each user interface and the set of data collectors in each data center.
    b. One VIP address for data collectors set to sticky session by IP for APM Gateway connection.
    c. One VIP address for APM user set to sticky session by IP for APM Gateway connection.
6. VIPs need to be configured for both local load balancers and for the wide area load balancer.
7. APM and SM would be integrated with the UCMDB.
8. Fail-over from site to site would be via a change in the wide area load balancer to force all traffic to the second site. *Note that users currently logged in would be disconnected and any work would be lost.*
9. All APM configuration (alerts, saved user reports, key performance indicators, health indicators, service level agreements) on either site needs to be manually replicated in the other site.

This architecture features two distinct APM databases which are not replicated. Manually duplicating the configuration (alerts, saved User Reports, key performance indicators, health indicators, service level agreements) and data collection essentially replicates the metrics, topology, and configuration.

Conclusion

The four types of redundant solutions described above (APM High Availability, APM Distributed High Availability, Disaster Recovery, and Business Continuity) are each possible with HPE APM9.30.

Each solution has its advantages and disadvantages. Each operations team using HPE APM needs to take a close look at the benefits and drawbacks of each solution, and decide which meets their needs. For some operations more than one solution will be needed. We can envision many cases where both APM High Availability, and either APM Disaster Recovery or APM Business Continuity solution - will be deployed.

Send Documentation Feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on APM Redundancy - Best Practices (Application Performance Management 9.40)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to docteam@hpe.com.

We appreciate your feedback!