BSM Redundancy

Prepared by:

HP Software

May, 2014

Date Prepared: 05/28/2014

Proprietary Notice

Revision History

Revision	Date	Comments
1.0	04/25/2014	Initial Revision
1.1	05/05/2014	After Initial Review
1.2	05/06/2014	After Internal Comments
1.3	05/08/2014	Added Distributed High Availability Option and made more corrections
1.4	05/09/2014	Additional Changes after Internal Comments
1.5	05/14/2014	Additional Changes to Business Continuity section
1.6	05/28/2014	Final Review

Table of Contents

Introduction

HP Business Service Management (BSM) 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 centre. This software can include:

- Applications Performance Management (APM) – software designed to keep your business healthy by monitoring applications across traditional, mobile, virtual and cloud environments. It 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.
- Operations Management I (OMi) - universal event-correlation software for diverse IT domains. Via the HP Run-time Service Model (RTSM), OMi uses IT topology to automatically correlate related events for quicker and easier root-cause identification—essential in today's complex virtualized and cloud environments—and for heightened efficiency of ITIL event and incident management.

In this paper we will look at four alternatives for BSM Redundancy. **BSM High Availability** and **BSM Disaster Recovery** have been well-documented. We will look at two additional alternatives: **BSM Distributed High Availability** and **BSM Business Continuity**. Each of these four use cases will be covered in detail in this document, this coverage will include architectural alternatives, advantages and disadvantages of each possibility and best practices.

There are two metrics we look closely 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 with HP BSM using the **BSM Business Continuity** use case. The process for doing this, however requires a considerable amount of customization and is often times not determined to be a viable alternative. The hope is that by reviewing this document each customer can choose the best solution for his data center.

It is also desirable to examine redundancy for BSM Data Collectors (Operations Manager, Network Node Manager, Business Process Monitor, Real User Monitor, BSM Connector, Diagnostics, Transaction Vision and SiteScope) as well. We will not cover solutions for data collectors in this document. It may also desirable to examine redundancy for the UCDMB. We will not cover solutions for UCMDB redundancy in this document.

BSM High Availability

Introduction

HP BSM can be set up in a High Availability configuration. Setting up High Availability for BSM 9.24 is documented in the BSM 9.24 Installation Guide – Appendix F: High Availability for BSM (referred to in this document as simply "Appendix F").

Setting up High Availability involves load balancing a set of BSM Gateway Servers and setting up a Failover BSM Data Processing Server. Load balancing the BSM 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

- BSM High Availability is a fully-supported out-of-the-box solution provided by Hewlett-Packard
- BSM 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. Metrics, events and topology coming in from data collectors however should be persisted and recoverable once the backup DPS is operational.
- Recovery Time Objective is lower than BSM Disaster Recovery

Disadvantages

- Recovery Time Objective is greater than zero. Realistically we can expect an RTO of 15 minutes – 60 minutes depending on your environment. This depends on how long it takes to start up the backup DPS processes.
- BSM 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. BSM 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.
- BSM High Availability acts as a single BSM instance rather than two independent BSM instances

Architecture

High Availability for BSM 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 load balanced by a local load balancer. A primary and failover Data Processing Server is deployed in the same data center.

There is only one logical database server utilized in BSM High Availability, so no synchronization of events, metrics and topology is needed. 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.

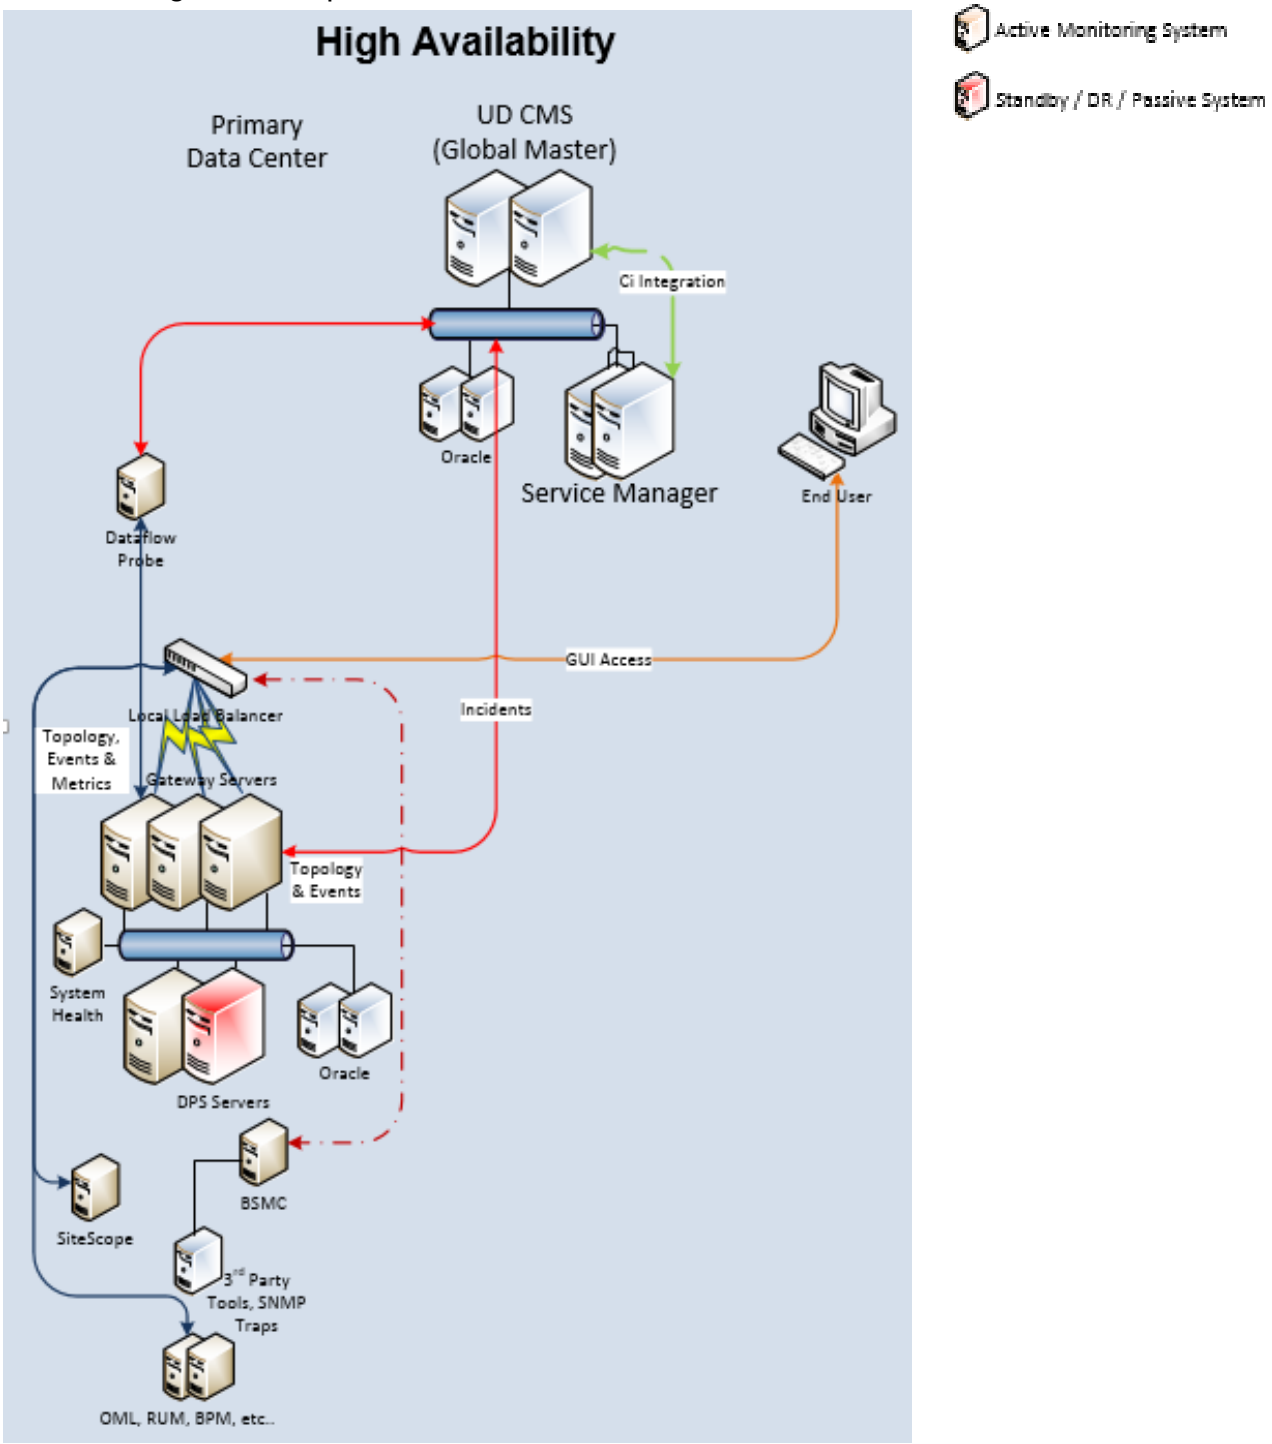

Best Practices

- Reference the BSM 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 BSM Data Processing server's services, follow the instructions in the **Configuring Automatic Failover** section of Appendix F (page 91).
- 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 BSM 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 BSM System Health. BSM System Health has a user interface that will allow for quick re-assignment of services.
- Automatic Data Processing Server Failover can be enabled by using the process in Appendix F (page 91), but an easier way to accomplish this is by using BSM System Health. BSM System Health has a user interface that will enable automatic failover:

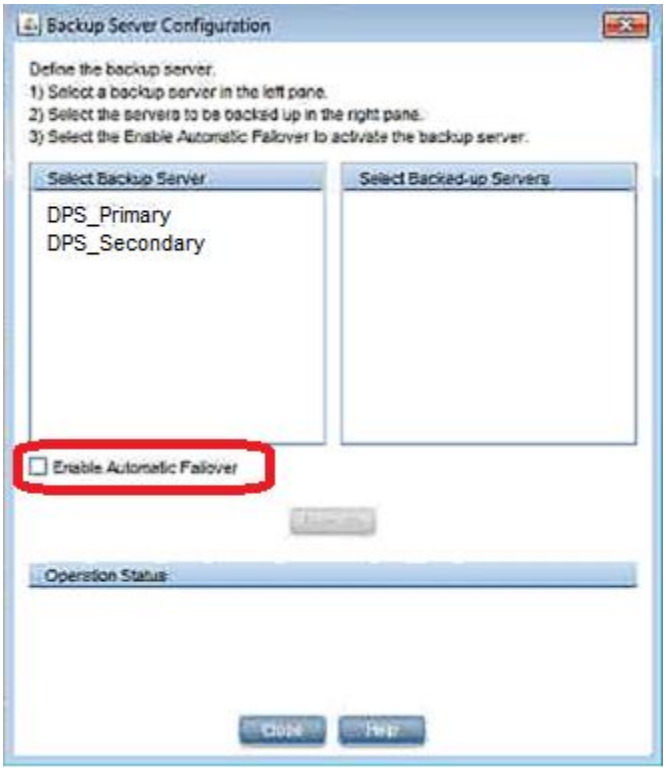

BSM Distributed High Availability

Introduction

HP BSM can be set up in a Distributed High Availability configuration. Setting up Distributed High Availability is similar to setting up BSM High Availability for BSM 9.24 documented in the BSM 9.24 Installation Guide – Appendix F: High Availability for BSM (referred to in this document as simply "Appendix F").

The key differences are:

- Deployment to distributed (two separate) data centers
- The requirement for a high speed network (<5ms latency) between the two data centers.
- The requirement for a distributed BSM database (see Disadvantages below)

Setting up Distributed High Availability involves load balancing two sets of BSM Gateway Servers (up to three servers in each set) and setting up a Primary and Failover BSM 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 BSM 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 process have been started.

There is only one logical database server utilized in BSM High Availability, so no synchronization of events, 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

- BSM Distributed High Availability is a fully-supported out-of-the-box solution provided by Hewlett-Packard
- BSM Distributed High Availability is a reasonably inexpensive solution to maintain
- BSM 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. Metrics, events and topology coming in from data collectors however should be persisted and recoverable once the backup DPS is operational. A manual process, however is needed to fail over which may impact your recovery point depending on how long it takes to execute this process.
- Recovery Time Objective is lower than BSM Disaster Recovery

Disadvantages

- Recovery Time Objective is greater than zero. Realistically we can expect an RTO of 30 minutes – 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.
- BSM Distributed High Availability acts as a single BSM instance rather than two independent BSM instances
- Requires a distributed BSM 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. This necessitates the two data centers must be within 400 km (250 miles) from each other with dedicated fiber for the BSM application between each data center (see Best Practices section below).

Architecture

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

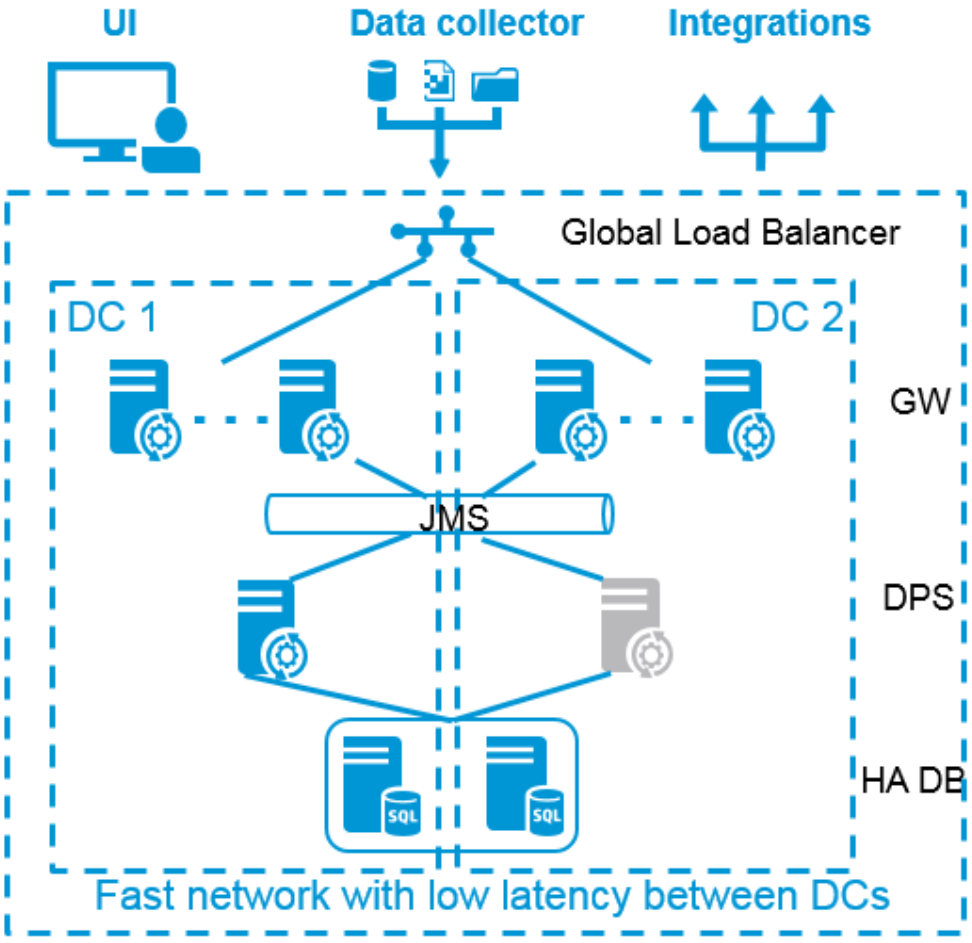

A more detailed implementation diagram is shown on the next page. This shows the BSM Gateway servers and BSM Data Processing Servers split across the two different data centers. A database cluster is also housed with nodes in each data center. The BSM Database acts as a single database clustered across 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 BSM using either:
  - The JMX Console (instructions shown in Appendix F, page 93 – Reassigning Services with JMX Console)
  - BSM System Health

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

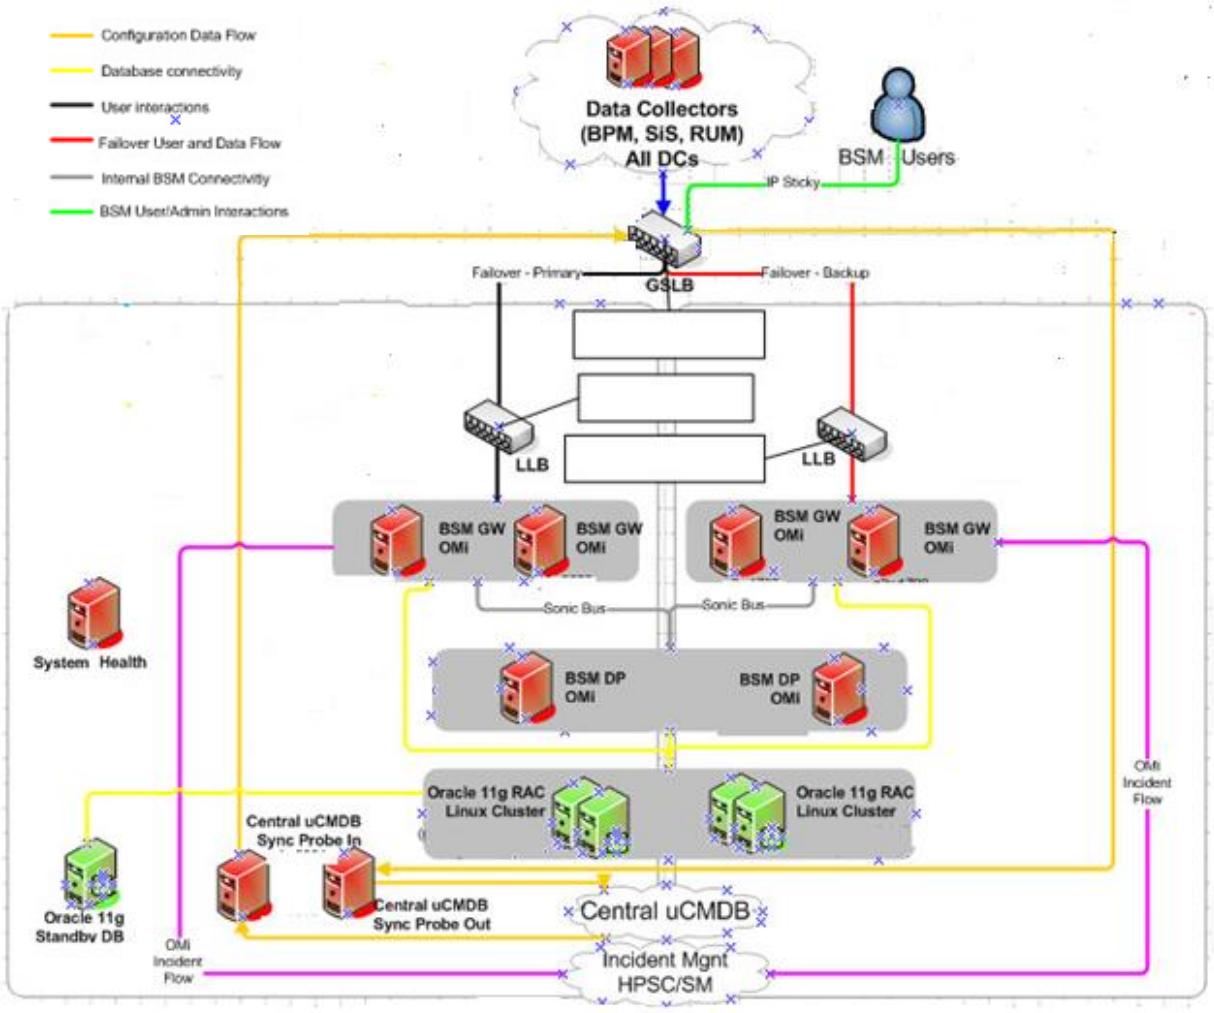

Best Practices

- Reference the BSM High Availability Fine Tuning Best Practices document
- Best Practices for this solution are identical to the Best Practices illustrated for BSM High Availability section of this document. There are two additional best practices:
- Ensure you have a consistent, <5ms connection between the two data centers. Please note:
    - 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 (Source: http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DCI/4-0/EMC/dciEmc/EMC_2.html , search for "**Based on the speed of light"**).
    - This would assume dedicated fiber capable of carrying the communication bandwidth for the BSM infrastructure between the two data centers.
    - It is important to benchmark the BSM application and BSM 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 BSM 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 stressful time. Practice will reduce your Recovery Time Objective.

BSM Disaster Recovery

Introduction

Duplicate HP BSM instances can be set up in a Disaster Recovery configuration. Setting up Disaster Recovery for BSM 9.24 is documented in the BSM 9.24 Installation Guide – Appendix E: Disaster Recovery for BSM (referred to in this document as simply "Appendix E").

Setting up Disaster Recovery involves setting up two complete BSM systems. The BSM Production Instance, with one to three Gateway Servers and one or two Data Processing Servers and the BSM Failover Instance identical to the BSM Production Instance (a second set of Gateway and Data Processing servers). In this configuration there are separate logical databases for the BSM Production instance and the BSM 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 BSM database schemas**. This procedure requires at least one BSM Administrator and one database administrator, who is familiar with the BSM databases and schemas.
- There are a number of different possible deployment and configurations for BSM. To validate that the disaster recovery scenario works in a particular environment, it should be thoroughly tested and documented. You should contact HP 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 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

- BSM Disaster Recovery is less expensive to maintain than BSM Business Continuity and still allows for recovery in a true disaster situation
- BSM Disaster Recovery is not restricted to a single data center
- There is not 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

- BSM Disaster Recovery is more expensive to maintain than either BSM High Availability or BSM Distributed High Availability
- Recovery Point Objective is highest of four 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.
- 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:

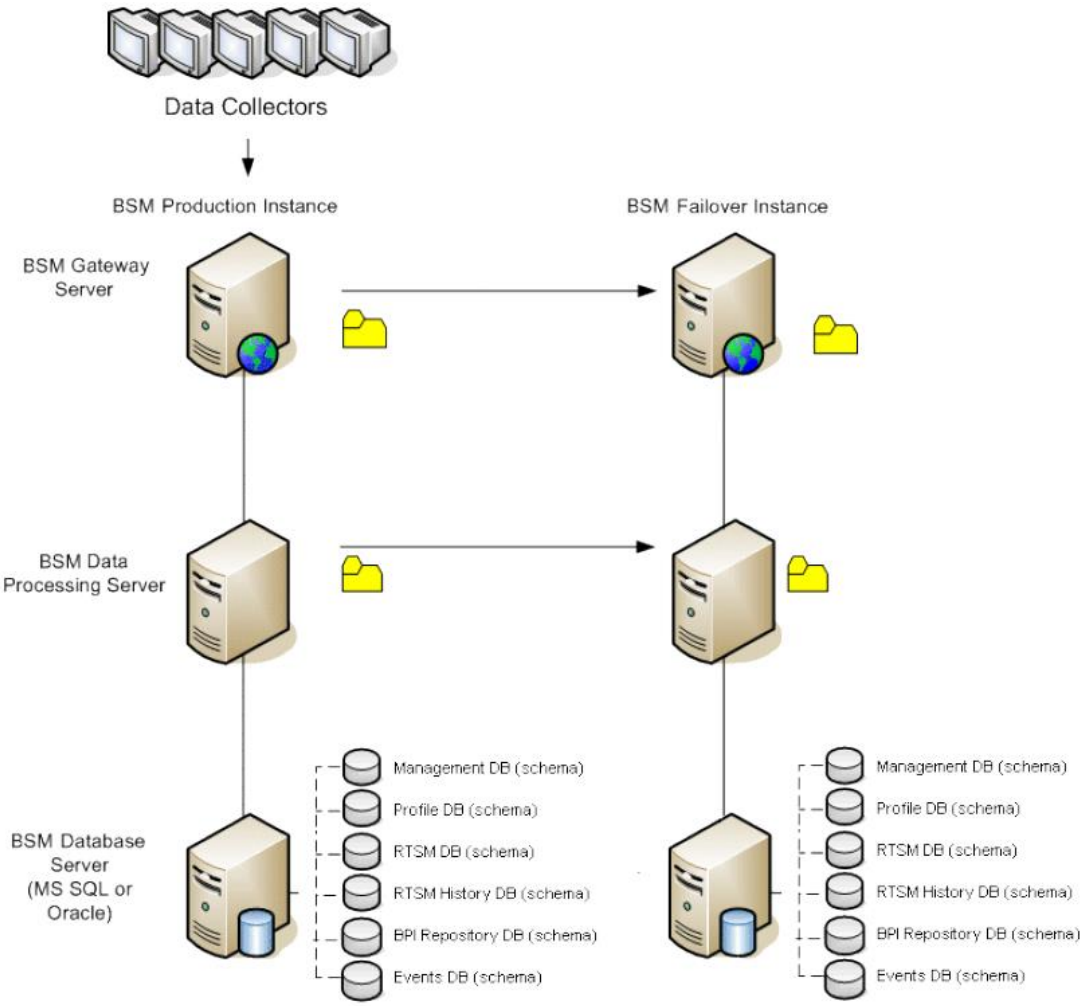

In the example above (with no external UCMDB) database replication is used to replicate **all** of the BSM databases. In the example on the next page, we are choosing to use an external UCMDB as the topology system of record. Therefore, we would only replicate the Management, Profile, BPI Repository and Event databases. Topology synchronization from the external UCMDB would keep the RTSM databases synchronized.

This diagram can be modified to accommodate for UCMDB synchronization. In our case we have data populating from the Universal CMDB's CMS Database to each of the BSM RTSM instances (both Production and Failover). This populates necessary CIs in the RTSM that are being discovered by Universal Discovery.

We have the Production BSM Instance populating CIs back to the Universal CMDB's CMS Database. This populates necessary CIs in the Universal CMDB that are being discovered by the BSM Data Collectors:

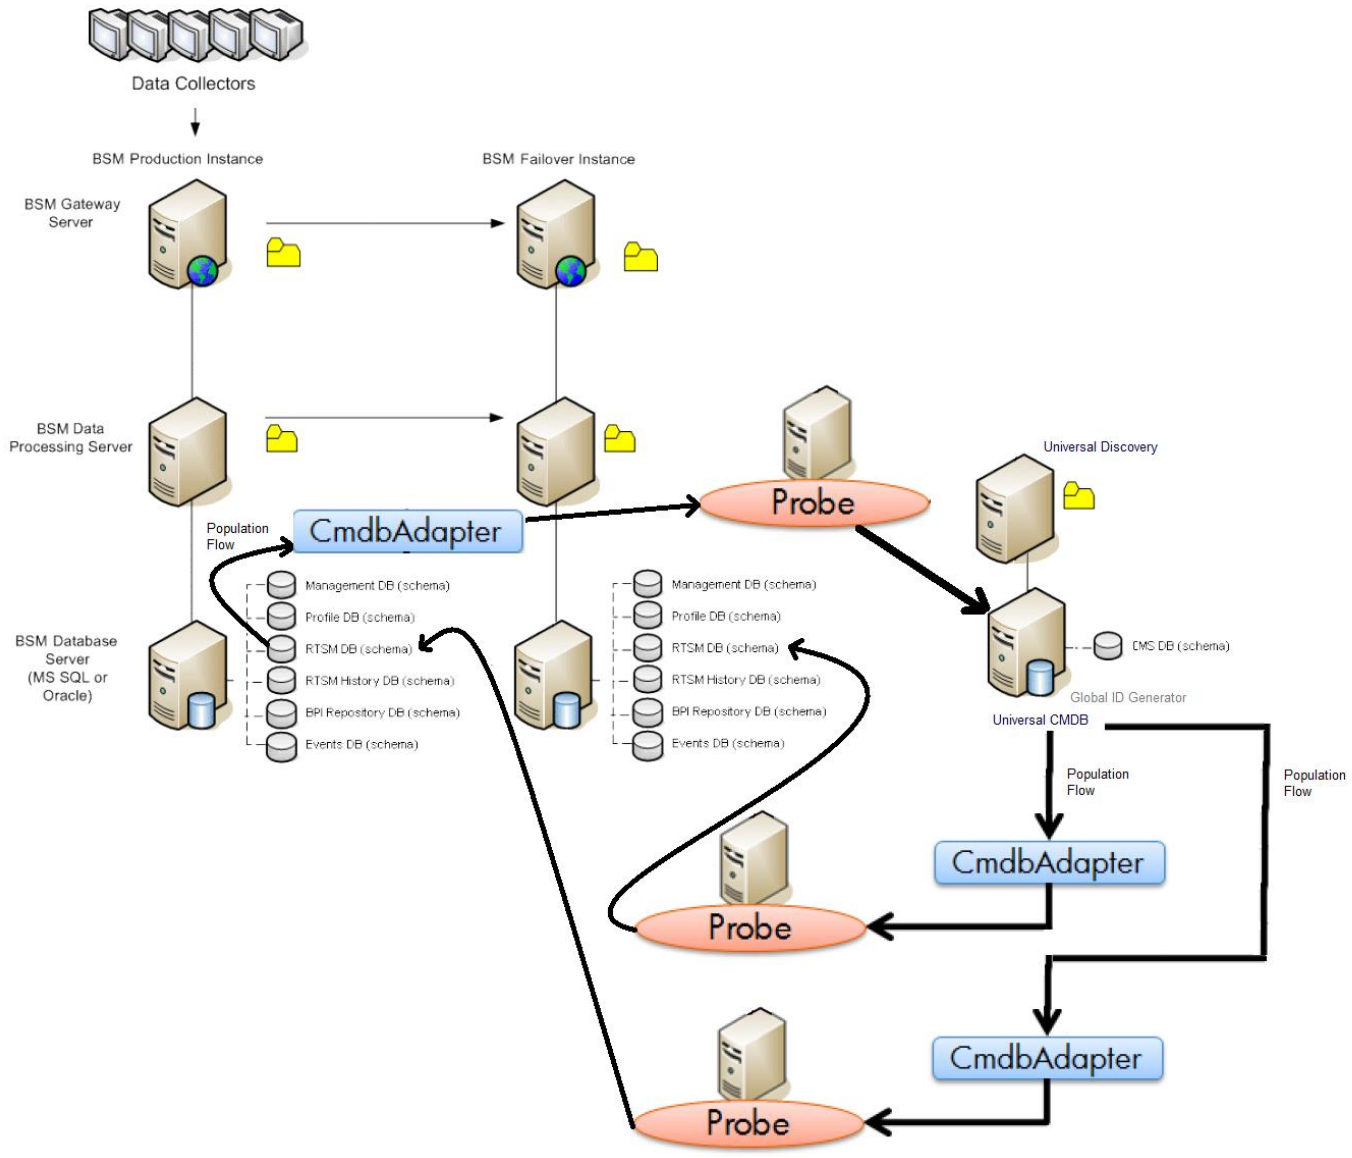

The environment in your data center may vary considerably from this, and accommodations may need to be made for:

- 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 Clean-up Procedure For Your Environment

It is important also to note that the procedure to failover is a very manual process. Appendix E shows a Clean-up Procedure section that is over five pages long and that 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:

- **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. This includes:
  - <context value>
  - <new value>
  - <key>
  - NewDatabasehostname
  - NEWDatabaseServerName
  - NEWSID
  - OLDSID
  - NEW_UID_name
  - OLD_UID_name
  - NEW_port_name
  - OLD_port_name

  These values will be known and a new document should be created with these values populated. If this exercise is completed prior to needing to do a Disaster Recovery confusion and mistakes can be avoided. Time can also be saved from needing to look these values up.
- **Practice:** If Disaster Recovery is practiced 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

The BSM 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 created challenges for even the most experienced administrators.

We recommend:

- A thorough review of Appendix E
- Modification of the diagrams in Appendix E to conform with your environment
- Modification of the Clean-up Procedure in Appendix E to conform with your environment
- Continued practice of disaster recovery during schedule outages and continual modification of your internal documentation to correct for changes in your data center

Using these practices can create an environment where unforeseen problems can be recovered from confidently.

BSM Business Continuity

Introduction

Duplicate HP BSM instances can be set up in a Business Continuity configuration. Setting up Business Continuity for BSM 9.24 is similar to setting up BSM Disaster Recovery (documented in the BSM 9.24 Installation Guide – Appendix E: Disaster Recovery for BSM).

The key differences are:

- The two BSM Database systems are separate and independent of each other.
- No data is replicated between the two BSM Databases
- All configuration, including data collector integration, alert configuration, downtime configuration, saved User Reports, RTSM Packages and OMi Content Packages must be manually be entered in on both systems

Setting up Business Continuity involves setting up two complete BSM systems. The BSM Production Instance, with one to three Gateway Servers and one or two Data Processing Servers and the BSM Failover Instance identical to the BSM Production Instance (a second set of Gateway and Data Processing servers). In this configuration there are separate logical databases for the BSM Production instance and the BSM Failover Instance. The two logical databases in this solution are not replicated.

Advantages

- BSM 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
- BSM Business Continuity can be used to reduce the user load on Gateway servers as users can be set to use the closest BSM system
- BSM Business Continuity can be used to reduce the data entry load on Gateway servers and Data Collectors. This may result in the ability to increase polling durations. For example, a single BSM instance with 1 SiteScope running all monitors at 5 minute intervals can be split up across 2 BSM instances with 2 SiteScope servers running all monitors at 10 minute intervals. This effectively provides 5 minute polling intervals with only half the load on each BSM instance.

Disadvantages

- BSM Business Continuity is by far, the most expensive solution to maintain as each data collector, each monitor and each configuration (alerts, saved User Reports, key performance indicators, health indicators, service level agreements) needs to be duplicated on each system
- BSM Business Continuity requires discipline in managing all of the BSM configuration (alerts, saved User Reports, key performance indicators, health indicators, service level agreements) on two separate BSM instances
- BSM Business Continuity requires duplication of most BSM data collectors (SiteScope, Business Process Monitor, Real User Monitor and BSM Connector). Duplications of some data collectors (Operations Manager) may not be required, but may be used to establish Business Continuity at the data collector level.

- Metrics and Events from data collector may not match each other precisely on both instances. They are not time-synchronized to the millisecond. For example, Business Process Monitor #1 connected to BSM #1 may report that a transaction failed. Several milliseconds later Business Process Monitor #2 connected to BSM #2 may report that the transaction succeeded. This may 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, one data collector that monitors a target system once every 5 minutes could turn into two data collectors that monitor a target system once every 5 minutes. This can be accommodated for by increasing intervals -- because we have two data collectors monitoring the target system, decision makers may agree to set the interval for each data collector to once every 10 minutes.
- BSM Business Continuity will double the number of licenses needed for BSM and each of its data collectors
- BSM Business Continuity has a customized integration between the two OMi instances. This requires the development of Groovy scripting to prevent duplicate incident creation to keep events synchronized between the two operations bridges within HP BSM.
- BSM Business Continuity is currently a customized integration between Service Manager, Operations Orchestration and HP BSM instances. This requires the development of Groovy scripting to ensure that duplicate incidents are not opened in Service Manager and that duplicate run books are not generated in Operations Orchestration.
- BSM 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 to the other. Duplicating configuration entries on two different systems is a process that requires a high amount of discipline and maturity for the team managing the solution.
- The BSM Business Continuity solution is not a fully tested, fully documented or fully supported approach.

Architecture

The architecture of BSM Business Continuity is illustrated below:

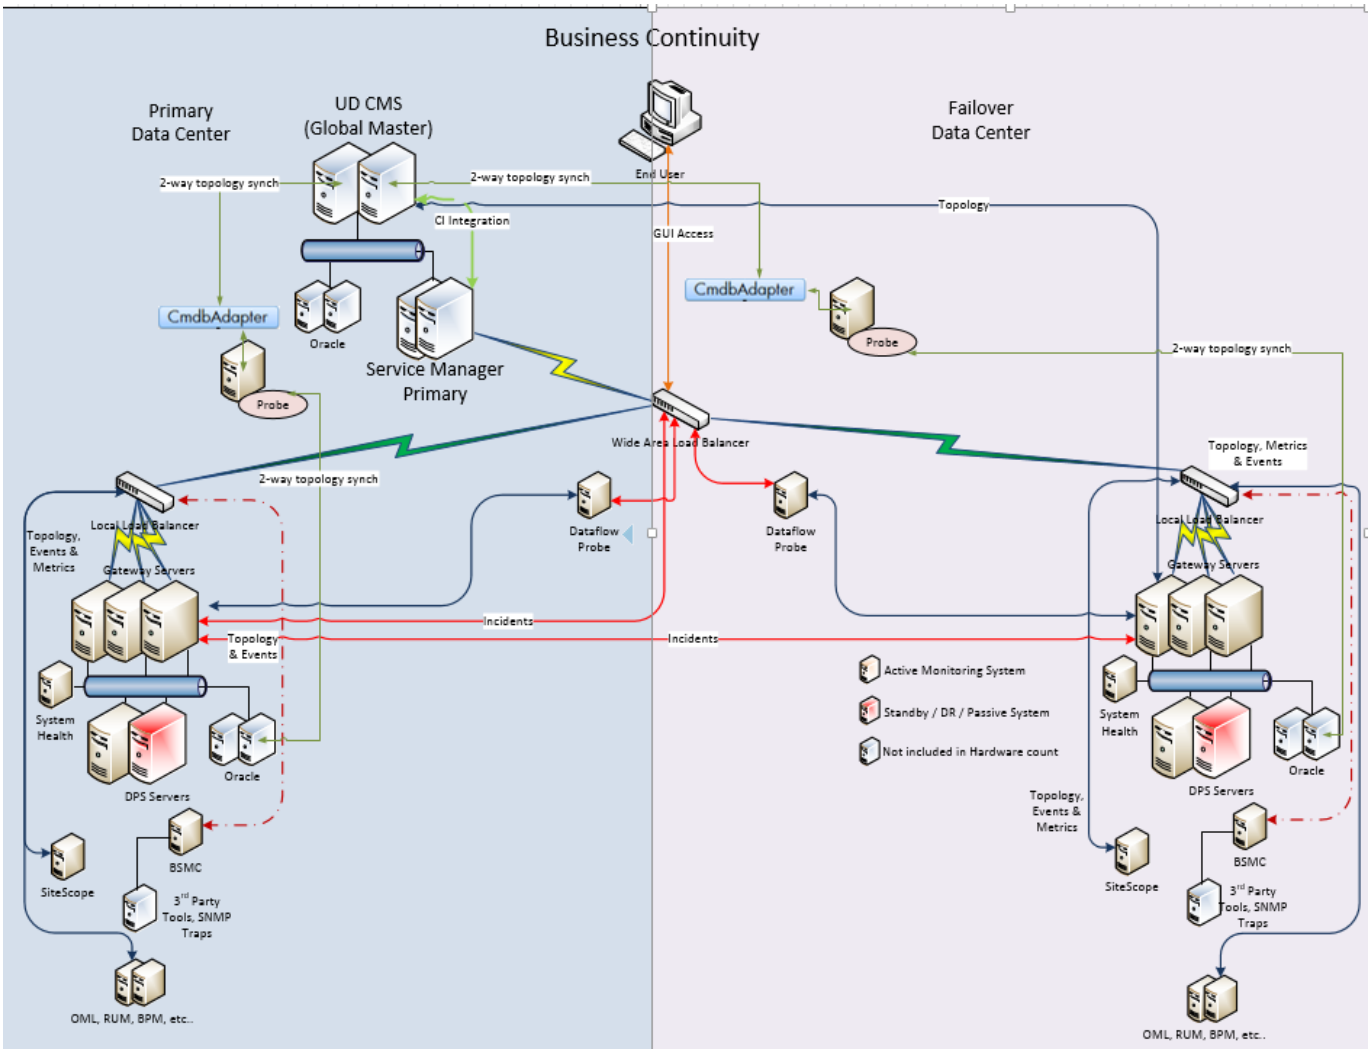

This architecture requires:

1. Two stand-alone implementations of BSM (one in each Data Center).
2. Integration with UCMDB (Master for Global ID's).
3. Integration of both BSM instances with Service Manager via Wide Area Load Balancer. This is currently a customized solution where the first BSM instance acquires the event marks the event as the owner, forwards it to the second BSM. The second BSM instance sees the same event, but only increments the counter and forwards the same event back to the first BSM 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 will require 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. BSM user VIP address would be the target (set to sticky for current Active BSM local user VIP)
5. Two Local area load balancers to point users and Data Collectors at the Local BSM's.
    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 BSM gateway connection
    c. One VIP address for BSM user set to **sticky session by IP** for BSM gateway connection
6. VIPs will need to be configured for both Local Load Balancers and for the Wide Area Load Balancer.
7. BSM 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 second site. *Users currently logged in would be disconnected and any work would be lost.*
9. All BSM 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 BSM 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 for us.

Conclusion

Four types of redundant solutions described above (BSM High Availability, BSM Distributed High Availability, Disaster Recovery and Business Continuity) are each possible with the current version of HP Business Service Management (9.2x).

Each solution has its advantages and disadvantages. Each operations team using HP Business Availability center needs to take a close examination of 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 with both a BSM High Availability and either a BSM Disaster Recovery or a BSM Business Continuity solution will be deployed.