



dynaLearn

Dynatrace HA installation in practice

May 18, 2016

Page is a Work in Progress

Several key components must be in place to accomplish HA for the DynaTrace server

As we take a deeper look into the possible solutions for HA, we should consider the following setup.

Setup

I used two (2) Linux servers with Dynatrace 6.1 installed on each server.

Each Dynatrace server was started and an active license applied to each server. Since the license is tied to the machine/host, I needed one for each Dynatrace server setup.

A shared mount was created using SAN storage, this will allow us to share the data between the two servers. I used the mount point `/data/dynatrace` but this can really be anything as long as it is shared between the two servers.

I also used a single PostgreSQL database. We will discuss DB considerations as we delve into the solution more.

This configuration will use a hot (active) and cold (passive) server, which would be the most applicable setup.



You would also want to use a VIP (virtual IP) or DNS entry that would allow you to divert the incoming traffic from the Active server to the Passive failover server.

As you can see with this setup, I have two DynaTrace servers pointing to the same data for the sessions and the Performance Warehouse.

After I have the servers initially configured, the Passive failover server will be stopped. You can not have two dynaTrace servers running at the same time, connected to the same data resources.

If the Server fails due to a hardware or software failure, the Collectors buffer data for a period of time, ranging from 30 seconds to a couple of minutes depending on load and Collector heap configuration. As a result, no PurePaths or measurements are lost if the Server is started again within this period.

Scenario

Thus, I can start my Passive failover server, have my VIP or DNS redirect the traffic from the collectors and be up and running in just a few minutes.



In this scenario, we are also assuming the Active (hot) server has experienced a catastrophic failure or crash and is not recoverable.

As you can see, with this setup, I can easily switch to the passive server. Also, If I can accomplish the switch over within an acceptable time frame, then the chance of data loss will be significantly reduced.

Performance Warehouse

When looking at the database setup for the Performance Warehouse, there are several options which come to mind. Most Enterprise databases support HA, Oracle for example will have several acceptable solutions. Take a look at this page for more details on Oracle HA, http://docs.oracle.com/cd/B28359_01/server.111/b28281/architectures.htm#g1012375

SQLServer will also provide HA capabilities, <https://msdn.microsoft.com/en-us/library/ms190202.aspx>

PostgreSQL, <http://www.postgresql.org/docs/8.3/static/high-availability.html>

Thus, I can look at database strategies that will allow for HA, load balancing or replication of the data.

Scripts

While doing the testing of this setup, I also developed bash scripts that would assist with placing the needed files on the shared mount point (SAN/NAS). The scripts will also create symlinks for all the files and folders needed on both the Active and Passive server. With the symlinks created for specific files such as profiles, I will now be able to make changes and have all the changes reflected in the Passive server once I start it up.

Download and open the zip file to see the scripts written in bash. There is a script for the Active server and one for the Passive server. It will move the files to the SAN/NAS and create the symlinks needed.

[dynaTraceHA.zip](#)

Please direct questions about your Dynatrace product or comments on this document to:

Dynatrace Customer Support

<https://community.dynatrace.com/community/display/SUPPORT>

Copyright © 2016 Dynatrace LLC. All rights reserved. Unpublished rights reserved under the Copyright Laws of the United States.

U.S. GOVERNMENT RIGHTS-Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in Dynatrace LLC license agreement and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (OCT 1988), FAR 12.212 (a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

This product contains confidential information and trade secrets of Dynatrace LLC. Disclosure is prohibited without the prior express written permission of Dynatrace LLC. Use of this product is subject to the terms and conditions of the user's License Agreement with Dynatrace LLC.

Documentation may only be reproduced by Licensee for internal use. The content of this document may not be altered, modified or changed without the express written consent of Dynatrace LLC. Dynatrace LLC may change the content specified herein at any time, with or without notice. All current Dynatrace LLC product documentation can be found at <https://community.dynatrace.com/community/display/APMDOC>.

Adobe® Reader® is a registered trademark of Adobe Systems Incorporated in the United States and/or other countries.

All other company and product names are trademarks or registered trademarks of their respective owners.

Document exported by Roman Spitzbart on 18-May-2016 07:34 UTC.

Dynatrace is the innovator behind the new generation of Application Performance Management. Our passion: helping customers, large and small, see their applications and digital channels through the lens of end users. Over 5,700 organizations use these insights to master complexity, gain operational agility, and grow revenue by delivering amazing user experiences.

