

Understanding BAR Challenges and Improving Recoverability of MySQL Databases with NetVault[®] Backup APM for MySQL

MySQL Server is the world's most popular open source database with more than 10 million active installations. Many large organizations utilize MySQL database in support of business-critical enterprise applications. Our research reveals, however, that majority of these companies are not backing up their data due to a number of factors. The Backup and Recovery (BAR) challenges that were uncovered significantly reduce the users' ability to recover from data loss or corruption. Building insight into these high-risk scenarios and understanding NetVault[®] Backup APM for MySQL helps the database administrators and IT Managers to be proactive and eliminate downtime while increasing their confidence.

WHITEPAPER

COPYRIGHT AND TRADEMARK INFORMATION

Copyright © 2006 BakBone Software, Inc. All rights reserved..

Unpublished – Rights reserved under copyright and patent laws of the United States.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

All applicable federal, state and local laws govern user access to this notice. All information available in this notice is subject to U.S. export control laws and may also be subject to the laws of the country where you reside.

Although BakBone has attempted to provide accurate information in this notice, BakBone assumes no responsibility for the accuracy of the information. BakBone may change the programs or products mentioned at any time without notice. Mention of non-BakBone products or services is for information purposes only and constitutes neither an endorsement nor a recommendation.

ALL INFORMATION PROVIDED IN THIS NOTICE IS PROVIDED "AS IS" WITH ALL FAULTS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED. BAKBONE AND ITS SUBSIDIARIES DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

TRADEMARKS

BakBone®, BakBone Software®, BakBone Logo®, NetVault®, VaultDR®, Replicator™, On Demand Replicator™, Constant HA Cluster toolkit, other logos, slogans, and marks (collectively "Trademarks") that appear on within this notice, third party websites and BakBone's products are the property of BakBone Software, Inc. all rights reserved. Nothing contained on BakBone's notice should be construed as granting any license or right to use any Trademark without the prior written permission from BakBone Software, Inc. or the party that owns the prospective Trademark. The absence, in this notice of a BakBone product, service name, slogan or logo does not constitute a waiver of BakBone Software Inc.'s trademark or other intellectual property rights relating to such name, slogan or logo.

Table of Contents

Introduction	4
Top MySQL Backup and Recovery Challenges	4
Scripting MySQL Backups Increases Recovery Risk	4
Manual MySQL Recovery Leads to Increased Downtime	5
Lack of Integration with Backup Devices Reduces Business Continuity	6
NetVault®: Backup Application Plugin Module for MySQL	6
Increases Confidence and Reduces Risk While Deploying MySQL	7
Speeds Up Restores to Reduce Downtime	8
Ensure Business Continuity with Automatic Integration of Backup Devices	8
Summary	9
About Integrated Data Protection	9
About BakBone Software	9

WHITEPAPER

WHITEPAPER

Introduction

In today's business environments, branch offices, business units, and merged or acquired operations will often have different IT environments, making unified data protection of the entire business either impossible or prohibitively costly, thus hampering the agility of the organization. Businesses can't afford the risks associated with any critical data loss – and require reliable and flexible solutions to fit these new environments. These same businesses today are worldwide, around the clock operations, where downtime means a large impact on business operations.

The MySQL database has become one of the world's most popular open source databases due mostly to its consistent high performance, high availability and ease of use. MySQL has become the database of choice for a new generation of applications built on the LAMP stack (Linux, Apache, MySQL, PHP/Perl/Python). As MySQL is selected for more and more mission critical applications, downtime will have a greater impact on the organization's critical service level agreements and customer satisfaction levels; therefore, the need for MySQL data protection solutions that provide proven business continuity is becoming critical for business success.

Despite this compelling need, research reveals that 25% of MySQL customers surveyed are putting their data at risk by not even backing up their databases, while 61% of DBAs write customized scripts to backup their databases.

This whitepaper discusses the top backup and recovery challenges facing MySQL environments and how NetVault: Backup's APM for MySQL addresses the needs of MySQL environments. BakBone's flexible backup and recovery options provide rapid recovery from data corruption while increasing confidence and reducing risk when deploying MySQL. This document is intended for the MySQL DBA who is looking to understand the proper ways of protecting their MySQL environment.

Top MySQL Backup and Recovery Challenges

Scripting MySQL Backups Increases Recovery Risk

Since it is virtually impossible to predict the type of failure that might occur; a MySQL backup policy must be flexible enough to fit a number of recovery scenarios, which include but are not limited to:

- Recovery up to the point right before a drop table or other damaging SQL statement occurred
- Recovery up to and after an erroneous SQL statement
- Recovery after a system failure or disaster such as hurricane or other natural disasters

MySQL currently provides a utility available for backup processes, called `mysqldump`. However, the automation of this program into repeatable and scheduled backup policies is not something currently provided by MySQL. After first reading the discussion on MySQL backups within the MySQL Reference Guide, one might believe that wrapping several lines of automation code around a `mysqldump` statement is a fast and easy solution – however, it can soon morph into a complex nightmare when all the possible recovery scenarios are considered.

Furthermore, when scripting an effective MySQL backup policy, the users must take into account the following questions:

- How do I ensure that I will get a consistent backup across individual storage engines?
- How do I ensure that I do not lose any committed transactions during the backup?
- What are the risks associated with purging the binary logs?
- What will happen during the restore process if I do not purge the binary logs at the correct step in the backup process?
- How do I implement a backup policy that will offer the most flexible restore options available?

Answering these questions requires a deeper understanding of the MySQL architecture and internal processes than can be found in standard MySQL backup reference guides. Thorough understanding of how these conditions impact the implementation of an effective MySQL backup policy requires a detailed investigation that entails its own share of trial and error testing.

Given the critical nature of managing a database environment, such as monitoring database health, performance tuning, and storage management, properly performing this investigation often takes a backseat. This increases the risk when a recovery is required because the existing backup sets will not be sufficient for the required recovery. In addition, the less experience the user has in recovering databases in a crisis situation, the greater the risk that they are going to miss critical steps in the process when scripting the “perfect backup policy.”

Manual Recovery of MySQL Leads to Increased Downtime

Scripting a MySQL recovery is not something that can be completed in advance due to a number of unknown factors, such as: determining exactly what needs to be recovered and at what point to perform the recovery. Meanwhile, incremental restores offer the greatest challenge when attempting to script a recovery because of the need to understand the exact order in which to apply the appropriate binary logs. (In order to preserve transaction integrity, the binary logs need to be applied in the order in which they were created.) Clearly, the answers to these questions cannot be predicted and writing a script in advance to handle all the possible scenarios only adds more complexity to the already complicated backup scripting.

Since scripting a recovery in advance is not a viable option, the MySQL restore commands would have to be issued manually. Issuing the commands manually, with the correct syntax in the correct order, is not an operation that anyone should be trying out for the first time they are trying to bring their systems back online as quickly as possible. Unless you have practiced these recoveries on a test system for multiple scenarios, the chances of getting the restore completed correctly the first time are very low. Unfortunately, if you perform the restore incorrectly, the entire process will have to be restarted. The longer this takes, the longer the system is unavailable.

Furthermore, manual MySQL restores require detailed MySQL knowledge and can not be performed successfully by someone other than an experienced MySQL administrator. Therefore, unless they are immediately available to start the restore process, the system will remain down until the DBA can be located and gain access to the system.

Lack of Integration with Backup Devices Reduces Business Continuity

MySQL's native utilities for backup and recovery, `mysqldump` and `mysqlbinlog`, do not provide integration with backup hardware devices such as tape libraries, standalone tape drives, virtual tape (disk) libraries, or shared virtual disk libraries. As important as having MySQL backups available online for faster restores, it is equally important to also store backups offline and offsite for business continuity purposes such as disaster recovery.

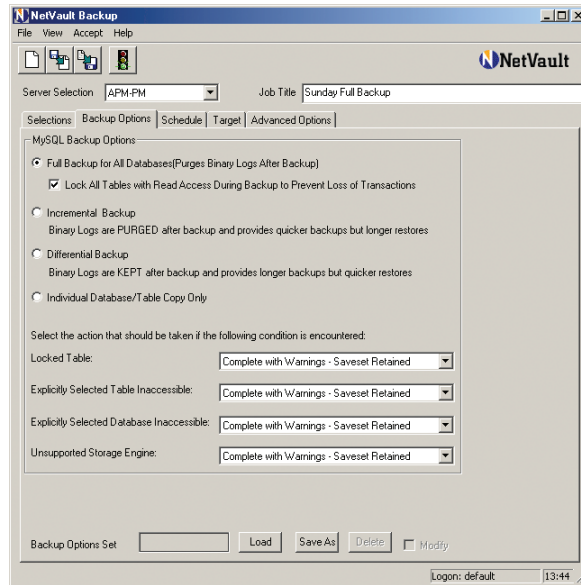
MySQL's lack of native integration with backup devices dictates that you must rely upon a Backup Administrator to ensure MySQL backups stored online are copied into offline storage at the correct interval. Furthermore, if an offline backup is needed for recovery, you must retrieve the appropriate backup set before starting the manual restore which results in an unnecessary increase in downtime.

NetVault[®]: Backup Application Plugin Module for MySQL

For data protection, disaster recovery and business continuity, the NetVault: Backup product line provides enterprise-class data protection for complex heterogeneous IT environments, regardless of their size. NetVault: Backup's flexible, modular architecture delivers proven reliability and high-performance, with an unmatched combination of features designed to meet current and future data protection needs. This unique flexibility reduces deployment costs and TCO by enabling customers to implement a NetVault: Backup solution that matches their current system and network resources. The NetVault: Backup optional components include NetVault: Backup Application Plugin Modules (APMs), NetVault: Backup Shared Virtual Disk, and NetVault: Backup SmartClients™.

“Price and performance are becoming more important as database workloads and operational requirements increase, and many organizations are looking to MySQL as a high performance, affordable database solution,” said Nicolas Pujol, director of Alliances, MySQL AB. “A fully-integrated enterprise data protection solution like BakBone's NetVault: Backup APM for MySQL can help ensure business continuity while enabling users to maintain the integrity of their data and applications more cost effectively.”

As a MySQL Enterprise Gold Certified Solution, the APM for MySQL extends the NetVault: Backup foundation product with specialized functionality designed specifically to address the top MySQL backup and recovery challenges.



NetVault: Backup APM for MySQL provides users with flexible and granular backup options to consolidate the backup and recovery of multiple MySQL storage engines into a single job, without the need for complex scripting.

create complex backup scripts and is flexible enough to account for a large number of recovery scenarios. Users no longer have to worry about understanding MySQL internals before being able to implement a backup policy that prevents committed transactions from being lost during backup and understanding the proper time to purge the binary logs. This knowledge has all been built into the NetVault: Backup APM for MySQL.

NetVault: Backup APM for MySQL's flexible backup features also include:

- Full, incremental, and differential backups while data is online and accessible
- Common UI across multiple storage engines
- Protection down to the table level
- Consolidation of multiple storage engines into a single job

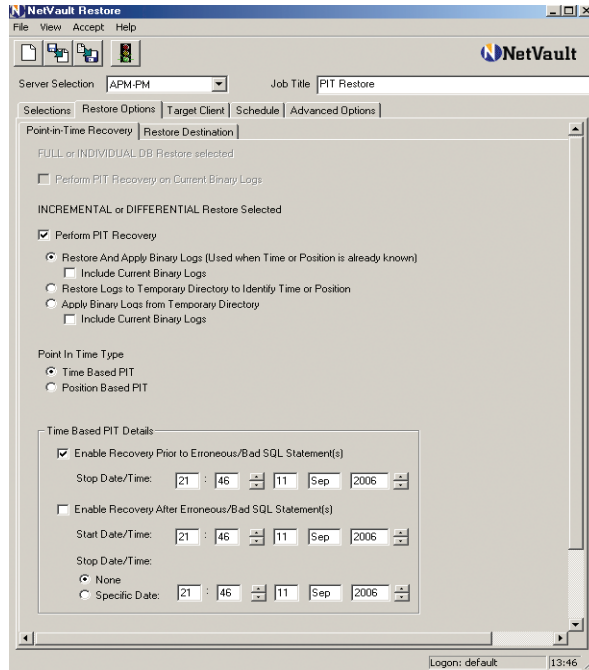
By relying on NetVault: Backup APM for MySQL to implement your backup policies, users can be freed to focus on more critical tasks without risking your ability to recover exactly what is needed in the unfortunate event of a failure. In addition, an IT managers' confidence is increased by knowing their MySQL data is protected, no matter what.

Utilizing an intuitive Graphical User Interface (GUI), NetVault: Backup APM for MySQL consolidates the backup and recovery of multiple MySQL storage engines into a single job without complex scripting. In addition to performing full backups or incremental, differential, or individual database while data is online and accessible, improved point-in-time functionality allows the NetVault: Backup APM for MySQL to perform more granular restores, letting users restore to a precise point in time and therefore significantly reducing data loss.

Increases Confidence and Reduces Risk While Deploying MySQL

NetVault: Backup APM for MySQL eliminates the need to

WHITEPAPER



NetVault: Backup APM for MySQL is integrated with MySQL's native point-in-time functionality, which allows for quick restores of MySQL environments and point-in-time recovery prior to and after data corruption occurrences.

Speeds Up Restores to Reduce Downtime

With NetVault: Backup APM for MySQL, you are no longer required to manually issue each command required in the restore process, nor are you required to determine exactly which binary logs need restoring and in what order they need to be applied. Simply select what needs to be restored, the backup set to restore from, and if appropriate, the time or position point to restore to and NetVault: Backup's APM for MySQL automatically performs the restore without further interaction. Not only are restores faster due to the reduced reliance on human interaction but the chance for a syntax error in a complicated script that would require a restart is eliminated.

Additional NetVault: Backup APM for MySQL restore features include:

- Full, incremental, and both time based and position based point-in-time restores
- Restores of complet instances, individual databases, or individual tables
- Rename of databases during restores
- Restores to alternate MySQL instances

Ensure Business Continuity with Automatic Integration of Backup Devices

With offsite backups being an important part of the data protection plan for any mission critical application, NetVault: Backup APM for MySQL takes advantage of NetVault: Backup's fundamental integration with a wide range of backup devices which include: tape libraries, stand-alone tape devices, virtual tape (disk) libraries, and shared virtual tape libraries. NetVault: Backup allows you the flexibility to select which backup device to store the backup on. You can store the backup online in a virtual tape library and duplicate the job to physical tape libraries shared by multiple MySQL instances, other propriety databases, or used for general backup purposes.

NetVault: Backup APM for MySQL gives you the confidence that your MySQL environments are being protected and stored offsite for disaster recovery purposes.

WHITEPAPER

WHITEPAPER

At the same time, it frees administrators from being available 24x7 because less experienced MySQL administrators can initiate restores with confidence that they will be performed correctly and as quickly as possible, thereby reducing downtime and improving business continuity.

Summary

When selecting MySQL as the database of choice for mission critical applications, confidence is increased with a MySQL Network Gold Certified Solution that addresses the top MySQL backup and recovery challenges. NetVault: Backup APM for MySQL provides proven business continuity for MySQL mission critical applications while offering a reliable disaster recovery solution. NetVault: Backup APM for MySQL also supports business growth and management of complex environments.

For more information about NetVault: Backup APM for MySQL, visit: http://www.bakbone.com/products/application_data_protection/mysql/

Download a free evaluation copy of NetVault: Backup APM for MySQL at: <http://www.bakbone.com/products/downloads/>

About Integrated Data Protection

BakBone's Integrated Data Protection™ (IDP™) strategy provides a comprehensive suite of integrated, heterogeneous cross-platform data protection technologies that are seamlessly managed from a single console. IDP enables administrators to easily and cost-effectively achieve high data availability, set schedules, manage policies, as well as control and report on their storage environment. BakBone's IDP framework consists of two key layers: the Core Services Managers™ and the individual data availability and protection technologies. BakBone's technology portfolio today consists of NetVault: Backup, a comprehensive backup and recovery solution, NetVault: Report Manager, advanced storage reporting and NetVault: Replicator, a replication solution for high availability.

About BakBone Software

BakBone Software is a leading international data protection solution provider that develops and distributes heterogeneous data backup, restore, disaster recovery, replication and storage reporting software for network storage and open-systems environments. Incorporated in 2000, BakBone Software's products are used by companies worldwide and are distributed through a network of OEMs and solution providers.