



DiskProphet™ BGP Case Study



Summary

Bureau of Geophysical Prospecting Inc., China National Petroleum Corporation(BGP/CNPC) has deployed ProphetStor DiskProphet in HPC (high performance computing) cluster system, to monitor disk status and predict disk failure. By integration with HPC task scheduler, potential failing nodes could be evaded in advance, right before task loading. The whole HPC system has been able to completely eliminate unplanned interruption, ensure that all tasks can be successfully accomplished in one time, thus shorten the average task running cycle more than 30%, improve overall computing resource efficiency by nearly 50%. The effective utilization rate of all out-of-warranty hardware, has been increased by more than 90%. The none-effective electric power consumption has been eliminated nearly 100%, and the operational workload has been significantly reduced.

The Company

BGP is the top data processing service provider worldwide, in oil and gas exploration industry, has occupied the biggest market share since 2015, in term of land exploration, BGP has won the biggest share for 15 years since 2002.

BGP's total assets are more than \$3.1 billion, including multiple data centers in China, United States, Middle East, and Southeast Asia. Based on ultra-high-density x86 HPC cluster and own developed GeoEast platform, BGP is providing exploration data analysis services to nearly 200 oil companies from 60 countries.

The Challenge

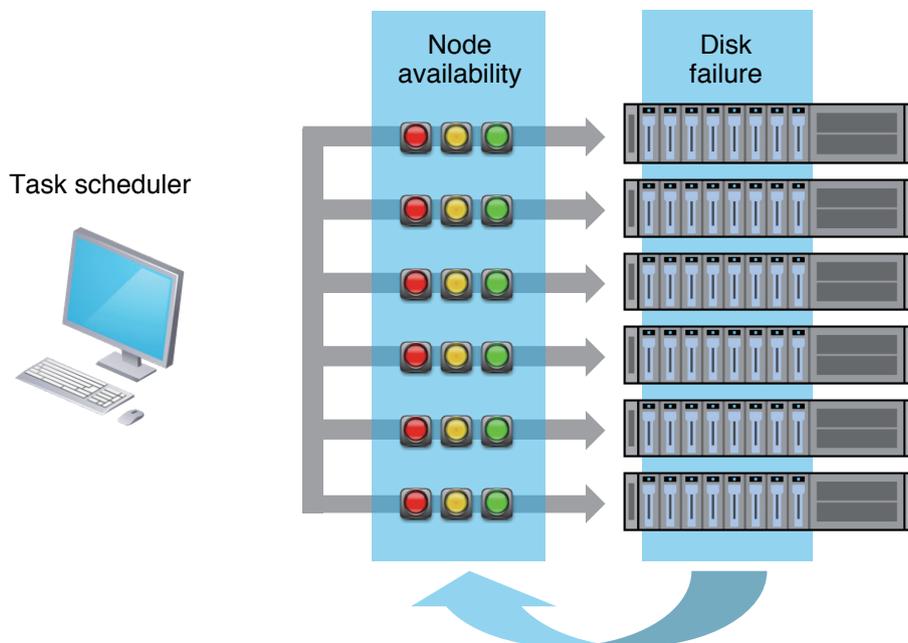
Exploration data analysis tasks are mostly massive and time consuming jobs, which need hundreds or thousands of nodes parallel processing for weeks or even longer. If unfortunately any unexpected node failure occurred during processing, it will inevitably result in reload of at least partial, if not all, task jobs. Therefore, even a tiny single node failure, would cause tremendous waste of time, computing resources, work force and energy consumption crossing entire cluster system.

Unfortunately, hardware failures especially disk failures are unavoidable in such large scale, high density, and heavily loaded clusters like BGP GeoEast system. So, how to smooth and stabilize the data analysis processing against disk failure, is the most important issue that has plagued the system construction and operation.

The Solution

Protection with redundancy, like RAID or multi-replica, is the most common way. But in actual environment of BGP, RAID parity or extra replica not only increase the hardware cost, but also significantly slow down the performance. After calculation, additional hardware investment plus extra processing time cost, has basically neutralized the transformation gains. So the traditional protection with redundancy, can not solve the problem, we must seek for new solutions.

With DiskProphet from ProphetStor, potential disk failures can be precisely predicted. By integration with GeoEast task scheduler, DiskProphet can help prevent jobs loaded on risky computing nodes, right before tasks start running. So that the entire cluster is guaranteed healthy enough through processing period of time.



After testing 1260 disks from 256 nodes, the accuracy of prediction within 45 days was up to 99%, accuracy level was ± 1 day. The prediction accuracy can fully meet real production requirements. BGP has now deployed DiskProphet onto 3000 production nodes in China processing center, and plans to expand its deployment to all processing centers around the world in near future.

Key Benefit Areas

- **Shorten data processing period of time**

With disk failure prediction by DiskProphet, the production task reloading is almost eliminated, and the average task execution time is reduced by more than 30%. As a result, BGP can provide the exploration data processing results much faster for oil companies without adding any hardware resources.

- **Leverage aged hardware resources**

In the past, because the node failure will affect the overall efficiency of the cluster, BGP used it cautiously, especially for old nodes that served over three years, only assigned for temporary or none-critical tasks. After DiskProphet deployment, the GeoEast scheduler can allocate tasks to any node, and the resources of all old equipment can be fully utilized.

- **Empower overall cluster system**

For the sake of proactive protection mode based on failure prediction, there is much less redundancy needed in HPC cluster. Even the most mission-critical tasks would just require fairly basic redundant configurations. In this way, a large number of resources are released, and can be used for effective production tasks. As a result, the overall resource efficiency has been greatly improved.

- **Simplify system management & maintenance**

Because of precise prediction, unexpected failures are reformed as planned events, and the handling procedure change from passive response to active governance. Therefore, the daily operation and maintenance work of the data center is stronger and the process management is more concise and efficient. The operation team can save a lot of time and attention from reliability, and focus on other work such as application optimization or new technology evaluation.

ProphetStor Data Services, Inc.

Headquarters

830 Hillview Court, Suite 100
Milpitas, CA 95035
☎ (408) 508-6255
🌐 <http://prophetstor.com/>

Taipei Office

4th Floor, No.86, Ming-Chuan Road
Shin-Dian District
New Taipei City, Taiwan 23141
☎ +886-2-8219-2814

Taichung Office

13th Floor, No.219, Minquan Road
West District
Taichung City, Taiwan 40341
☎ +886-4-2305-1816

China Office

#522, 5/F South Tower, Building C,
Raycom InfoTech Park,
No. 2 Kexueyuan South Road,
Haidian District, Beijing 100190
☎ +86 (10) 59822185

EMEA Office

ProphetStor Data Services SAS
2 place de Touraine
78000 Versailles, France
☎ +33 (0)1 70 29 08 66



PROPHETSTOR

Visit us at www.prophetstor.com
to find out more, email us at
info@prophetstor.com or contact
your local ProphetStor office.

Hong Kong Office

Unit 706, Haleson Building,
No. 1 Jubilee Street,
Central Hong Kong
Hong Kong, China

Malaysia Office

Level 28-03-03A, PJ Exchange
16A, Persiaran Barat, Petaling
Jaya Selangor, Malaysia 46050
☎ +60-3-79620108

Singapore Office

1 Raffles Place
Level 20-61 Tower 2, One Raffles Place
Singapore 048616
☎ +65-68085635

Tokyo Office

8F PMO Nihonbashi Mitsukoshimae,
3-4-5 Nihonbashi-honcho, Chuo-ku,
Tokyo 103-0023, Japan
☎ +81-3-6262-3936

Copyright © 2018 ProphetStor Data Services. All rights reserved.
ProphetStor Data Services and DiskProphet are trademarks or
registered trademarks of ProphetStor Data Services, Inc. in Taiwan
and other countries. All other company and product names
contained herein are or may be trademarks of the respective holder.