

# DiskProphet® BGP Case Study



## Summary

China National Petroleum Corporation (CNPC) deployed ProphetStor DiskProphet to monitor disk operational status, predict potential disk failures, and receive proactive alerts on soon-to-fail disks. The high-performance computing (HPC) center solved frequent disk failures that caused random host downtime during geophysical data analysis, which creates management overhead and results in a lower quality service. With DiskProphet, CNPC was able to conduct data analysis in the planned time. This proactive approach not only saves costs in both hardware and personnel costs, but also enables CNPC to provide better quality data analysis to customers.

## The Company

CNPC is one of the largest energy companies in the world. Its headquarters is in Beijing. CNPC ranked third in the 2016 Fortune Global 500. It was founded in 1988 with over one million employees. Its annual revenue is over \$240 billion. As a global energy company, CNPC collects and analyzes geophysical data when exploring oil all over the world. BGP, Inc., the subsidiary of CNPC, is responsible for operating a high-performance computing (HPC) center with over 3,000 servers to conduct geophysical data analysis.

## The Problem

In the HPC center, BGP/CNPC equips each host with directly attached disks (DAS). The DAS stores temporary data during computation. To speed up the data read/write I/O, the disks installed in each host are configured as RAID-0. Therefore, CNPC can combine the I/O of all host disks.

Due to their intensive data access during computation, the disks have higher failure rates than usual workloads. Disk failure is a severe issue for BGP/CNPC, because any disk failures give rise to host downtime. On some occasions, it even terminates the whole analysis process. Since data processing is conducted in parallel, having more available hosts to join the parallel data processing is critical.

## The Solution

BGP/CNPC can evaluate two options. The first is to purchase additional servers to fill the gap in computation resource loss due to disk failure. Soon, BGP will realize that adopting DiskProphet is the most cost-effective solution. BGP/CNPC can then replace the soon-to-fail disks that are predicted by DiskProphet. Another benefit of DiskProphet is its ability to apply historical performance metrics of any host and disk. The HPC center administrator can then use the time machine of the data to track unusual performance patterns at any point in time.

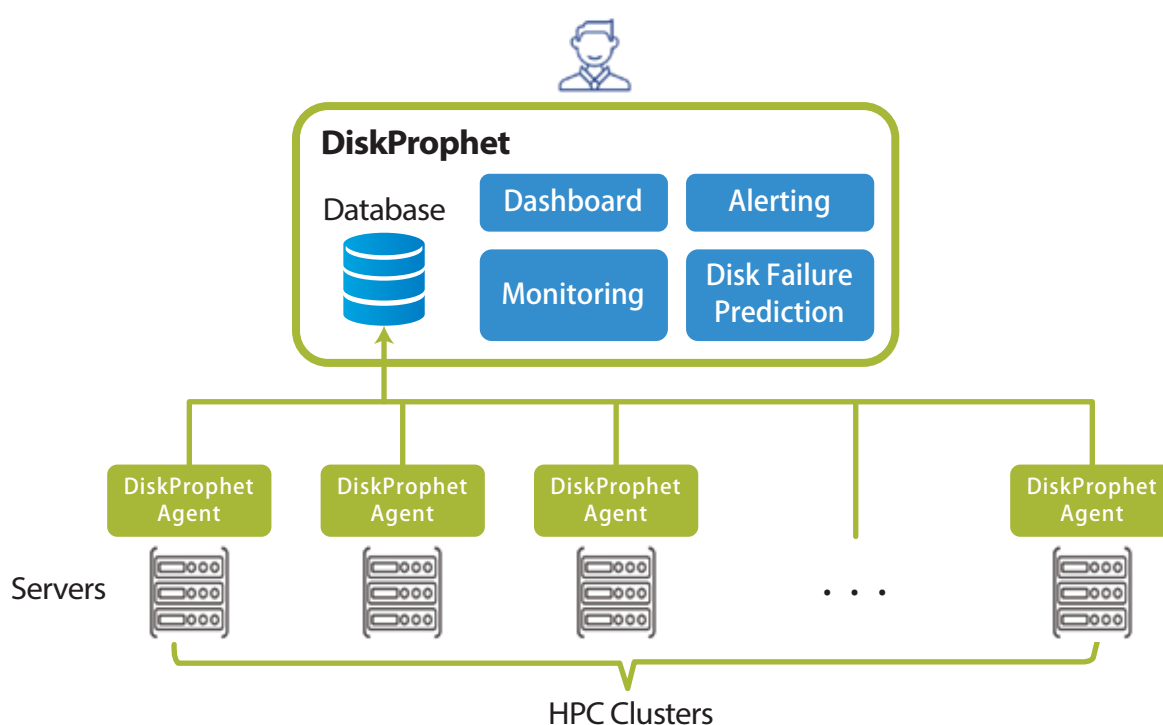


Figure 1. DiskProphet Architecture for BGP/CNPC

BGP/CNPC started the evaluation of DiskProphet by deploying a DiskProphet Agent to a 117-node cluster. The deployment process is quite simple because DiskProphet automates agent deployment to each host through Ansible. After thirty days of evaluations, BGP/CNPC was very satisfied with the benefits of DiskProphet brings. BGP/CNPC then turned to production deployment for the rest of 3,000 hosts.

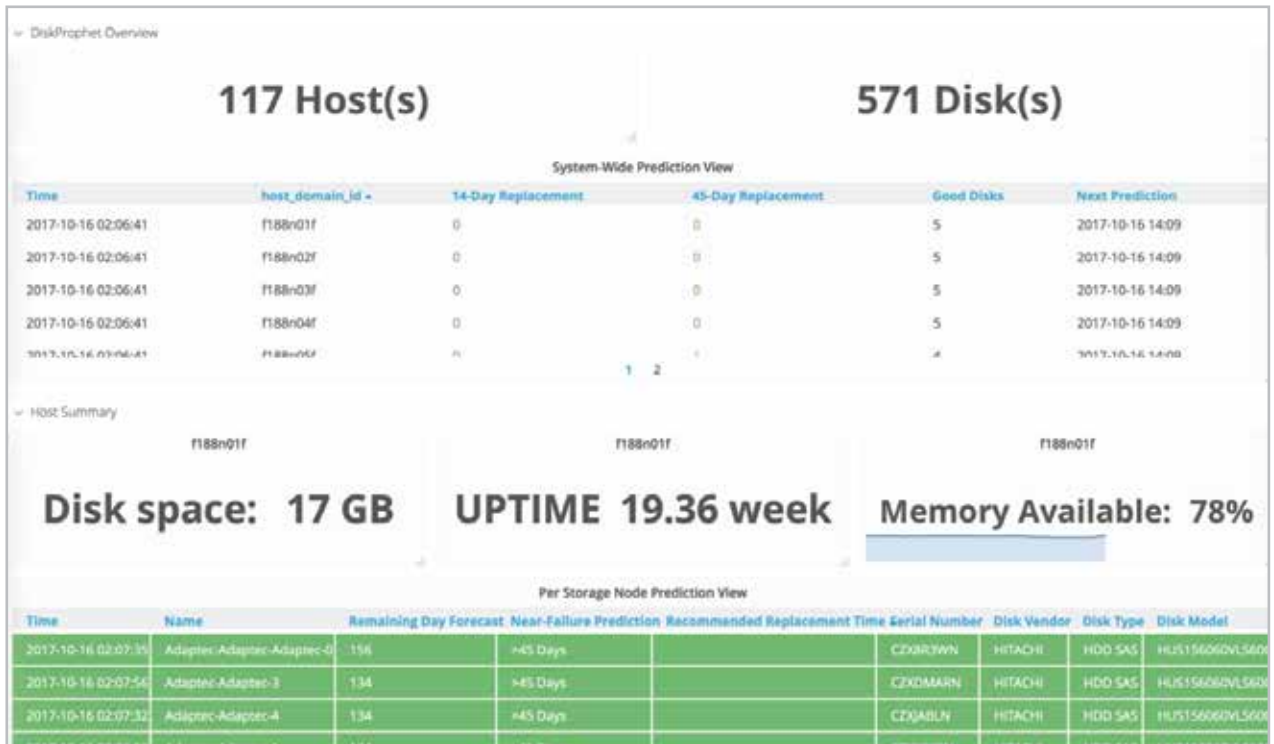


Figure 2. DiskProphet Dashboard Screenshot

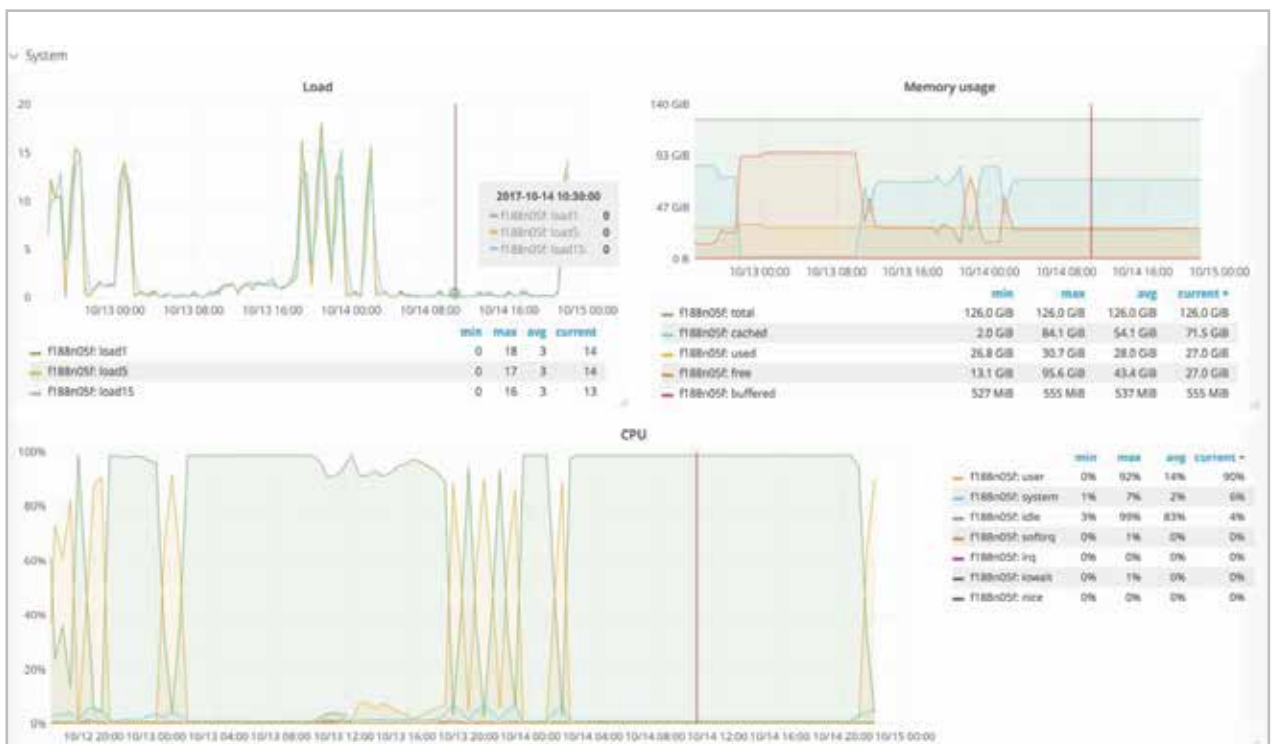


Figure 3. Using DiskProphet Dashboard for host monitoring

## Key Benefit Areas

Most of the DiskProphet benefits are to the HPC center operational cost.

- **Reduce Power Consumption:** Before deploying DiskProphet, nearly 1/4 of data analysis jobs needed to be redone due to disk failure. With DiskProphet, this is eliminated.
- **Avoid Engineer Hires:** With DiskProphet's proactive alerts on disk failure, BGP/CNPC turns random disk failure replacements to scheduled disk replacements. BGP/CNPC no longer plans to hire more engineers for hardware maintenance.
- **Improve Analytics Productivity:** BGP/CNPC researchers benefit from nonstop data analysis in the HPC center. Prior to DiskProphet deployment, they worried that the jobs wouldn't be completed in time, or that they might even be terminated prematurely.
- **Improve Sales Cycle:** Another important business to BGP/CNPC is to sell data analytics services to other energy companies in the world. BGP/CNPC finds that its salesforce can provide more efficient analysis services for winning energy exploration deals.

## ROI Calculation

Before the deployment of DiskProphet, BGP/CNPC purchased more hosts than its data analysis jobs needed. This is because over-provisioning to HPC clusters was its strategy to overcome disk failure. This leads to higher costs in many areas, including personnel, hardware, and power consumption.

It is standard for data centers and enterprises to acquire higher graded servers, or more servers than necessary to avoid SLA performance violations. During the DiskProphet evaluation process, BGP/CNPC not only assessed the usability and stability of DiskProphet, but it also created three-year financial protection when comparing DiskProphet deployment with the traditional over-provisioning approach. They found that the most significant portion of the ROI is the server procurement cut and the engineers' time spent on random disk failure replacements. Other than the aforementioned direct benefits, the indirect benefits included revenue gains from the higher efficiency data analysis service and much less interrupted computing process.

## Key Cost Area

The cost of deployment is mainly included software, hardware, and personnel.

- **Software:** Costs of DiskProphet include the initial subscription costs and licensing fees based on the number of disks to be monitored.
- **Hardware:** Initially, BGP/CNPC deployed DiskProphet on its existing virtual environment. They plan to add more hardware when they expand their HPC clusters.
- **Personnel:** DiskProphet is a self-driven software system which applies ProphetStor, the U.S.-patented technology. It is failover, scale-out, and self-healing. The BGP/CNPC maintenance team only accepted one day of training from ProphetStor, and they can operate DiskProphet.

### 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

2 place de Touraine  
78000 Versailles  
France  
+33 (0)1 70 29 08 66

#### 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  
#20-61 Tower 2, One Raffles Place  
Singapore 048616  
+65-68085635

#### Tokyo Office

〒103-0023 東京都中央区  
日本橋本町3丁目4-5  
日本橋三越前PMOビル8F  
+81-3-5542-1441



PROPHETSTOR

Visit us at [www.prophetstor.com](http://www.prophetstor.com)  
to find out more, email us at  
[info@prophetstor.com](mailto:info@prophetstor.com) or contact  
your local ProphetStor office.

Copyright © 2017 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.