



REDEFINING SOFTWARE-DEFINED STORAGE

How Lenovo and SUSE integrate leading hardware and software for best-in-class CEPH-based HPC storage

```
elif operation == "MIRROR_Z":
    mirror_mod.use_x = True
    mirror_mod.use_y = False
    mirror_mod.use_z = True
elif operation == "MIRROR_Z":
    mirror_mod.use_x = False
    mirror_mod.use_y = False
    mirror_mod.use_z = True

#selection at the back the deselected mirror
mirror_ob.select
modifier_ob = 1
bpy.context.scene.objects.active = modifier
print "Selected" + str(modifier_ob) # modifier ob is the active
mirror_ob.select = 0
name = bpy.context.selected
bpy.data.objects[name].name = "Mirror"
bpy.data.objects[name].name = "Mirror"
```

IDC estimates that the world's data is growing by 40% each year, ballooning to 44 zettabytes (i.e. approx. 44 million petabytes or 44 billion terabytes) in 2020 and 180 zettabytes five years after that.¹ This digital deluge will affect every enterprise and every HPC deployment. Those who stay on top of their rapidly growing data stand to reap rewards rather than lose out to their savvy competitors. To keep pace, software-defined storage is increasingly the most compelling solution to stand alone or — in larger HPC clusters — work in tandem with compatible primary storage options.

Lenovo and SUSE, powered by Intel® Xeon® Scalable processors, offer a portfolio of storage possibilities built around Ceph, a leading open-source, software-defined HPC storage solution.

SUSE Enterprise Storage, based on the open-source Ceph project, is a self-healing storage system that provides block, file and object data access in a single cluster, simplifying administration and reducing operational costs. Ceph's multisite object replication and asynchronous block mirroring provides uninterrupted serviceability, while its open source management framework provides a simplified interface and improved cost efficiency.

According to a 2018 customer survey, Ceph's users are most drawn to its open source nature and its extreme scalability. They also report strong customer satisfaction — more than 83% are either satisfied or "extremely satisfied" with the technology.²

¹ Saurabh Gupta, Venkata Giri, Practical Enterprise Data Lake Insights, Apress (2018) p. 5; "Double-Digit Growth Forecast for the Worldwide Big Data and Business Analytics Market Through 2020 Led by Banking and Manufacturing Investments, According to IDC," International Data Corporation (IDC), Oct. 2016

² "Ceph User Survey 2018 results," The Ceph Blog (ceph.com/ceph-blog), July 17, 2018

Yet, open source software's deep and wide and continuously updated code base can be a double-edged sword. In theory, any user can download Ceph and deploy it themselves. In practice, however, open source software that is self-deployed (and thus unmonitored and unmanaged) can sometimes offer a different user experience than what is possible "in theory."

After all, even Ceph's own customer survey — which includes those users who download and install Ceph on their own — reveals a further 17% who are not fully satisfied with their storage solution. For something so fundamental to the enterprise as storage, do-it-yourself deployment is long on risk and potentially short on reward.

SUSE has a long history of providing professional open source support services for enterprise customers. SUSE was founded in 1992, and SUSE® Linux is one of the oldest and most enterprise-focused Linux distributions in the industry. SUSE brings its substantial history of enterprise support to the Ceph space. SUSE supports its Ceph deployments with full service agreements, certification of hardware products and integration with the Linux stack and network environments.

Lenovo, with 10,000+ IT support specialists, 3,000 field service locations and 51 contact centers, offers a customer-focused infrastructure. Its four HPC and AI Innovation Centers (in Research Triangle Park, North Carolina; Stuttgart, Germany; Beijing, PRC; and Taipei, Taiwan) allow customers to optimize their applications, test at scale and validate their HPC solutions — including their SUSE storage solutions.

Use Case 1: Larger HPC clusters (200+ nodes)

Customers who run larger clusters consisting of more than 200 nodes typically require multi-tier storage solutions. Lenovo's preferred primary storage for larger clusters is its Distributed Storage Solution for IBM Spectrum Scale (DSS-G).

Spectrum Scale provides a global namespace and shared file system access among Spectrum Scale clusters. Additional features of Spectrum Scale include simultaneous file access from multiple nodes, high recoverability and data availability through replication, the ability to make changes while a file system is mounted, and simplified administration even in large environments.

DSS-G in this case, then, represents the cluster's Tier 1, high-performance data storage system.

For Tier 2 or Tier 3 storage (active archive and backups), historically tape storage has been the default. While cost effective, this solution has performance limitations. Data sent to tape is often "forgotten" and when needed retrieval times can be long, affecting enterprise productivity. Active archival solutions on the other hand give customers an option of further accessing the data with other analytics tools. Increased cluster speeds, general data growth and an explosion of unstructured data have all pushed enterprise IT — especially enterprise HPC — away from tape.

IBM Spectrum Scale feature called Transparent Cloud Tiering opens up a new opportunity to use scalable and cost-effective SUSE Enterprise Storage as the Tier 2 or Tier 3 storage component.

Consider, instead, the trusted partnership of Lenovo and SUSE deploying SUSE Enterprise Storage in a custom-configured and expertly deployed Ceph solution:

- Lenovo and SUSE solution and architecture support leverages extensive experience in SUSE Enterprise Storage systems deployments as well as testing and benchmarking some of the largest Ceph-based configurations.
- User enjoys a single point of contact for support — rather than a user being left to debug and coordinate with server, storage, OS and/or networking vendors themselves
- Lenovo and SUSE's hardware and software portfolio uses few, well-tested building blocks — rather than customer-installed, trial-and-error approach to Ceph deployment.





Transparent Cloud Tiering automatically migrates cooler or cold data to lower cost storage options, freeing up storage capacity for true Tier 1, high-performance needs.

SUSE Enterprise Storage with its socket-based pricing model is therefore an ideal cost effective solution to leverage Spectrum Scale features.

Use Case 2: Smaller HPC clusters (200 or fewer nodes)

Ceph is also a robust HPC storage solution which for a majority of Ceph users, surveyed in 2018, operate a raw capacity of 100 terabytes to 10 petabytes.³ In other words, Ceph represents robust Tier 1 as well as Tier 2 and 3 storage for the smaller HPC cluster operator.

Because Ceph is self-healing and self-managing, it automatically detects changing conditions and redistributes workloads if a node goes down. Built-in fault tolerance keeps the system working despite component failure and helps ensure that you don't lose any data.

SUSE Enterprise Storage exclusive node-based pricing model (as opposed to per-gigabyte models of other distributions) makes it the lowest-priced option for enterprise disk-to-disk backup. That means you don't pay more for saving more data.

Lenovo servers appropriate for a SUSE Enterprise Storage cluster include:

- Object storage device node — the devices in the cluster that do the actual data storage. Lenovo's ThinkSystem SR550 or SR650 (minimum 4 devices) is recommended
- Monitor node — maintaining information about cluster health and keeping maps of the cluster. Lenovo's ThinkSystem SR530 or SR630 (minimum 3 devices) is recommended.
- Gateway node — the translation between your backup server and the SUSE Enterprise Storage cluster. Lenovo's ThinkSystem SR530 or SR630 for small setups (adding machines as you grow) is recommended.

Conclusion

Lenovo and SUSE's Ceph-based storage is scalable, reliable and fail-safe. It runs on Intel® Xeon® Scalable hardware and has the ability to handle exabytes of storage. It is also open-source, freeing the user from future vendor lock-in that other, more proprietary solutions may present. Whether working in tandem with Lenovo DSS-G storage in a large HPC cluster or as a Tier 1 and Tier 2 storage solution in a smaller cluster, Ceph provides the scale and versatility that HPC operating environments increasingly demand. Consult your Lenovo partner or storage specialist today to discover how Lenovo, SUSE and Ceph can render your HPC storage future-facing and ready for the data needs of tomorrow.



CONTACT YOUR LENOVO REPRESENTATIVE

³ ibid., "Total raw capacity: From 100 TB to 1 PB: 37.84%, from 1 PB to 10 PB: 20.91%"

© Lenovo 2019. Lenovo, the Lenovo logo, System x, ThinkServer, ThinkSystem, ThinkAgile are trademarks or registered trademarks of Lenovo. company products and service names may be trademarks or service marks of others.

Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries.

© 2019 SUSE LLC. All Rights Reserved. SUSE and the SUSE logo are registered trademarks of SUSE LLC in the United States and other countries. third-party trademarks are the property of their respective owners.

