

# How SUSE Enterprise Storage Can Help You Manage Your Data Explosion

---

---

---

Whether you work for a large multinational corporation or an aspiring startup, chances are you're creating data at an alarming clip. Recent studies show that some 90 percent of all the recent internet data was created in the last two years. This includes startling numbers such as users uploading 400 hours of video to YouTube every minute, Instagram users sharing more than 46 million posts a minute and Twitter users posting about 12 million tweets every hour.

Of course, these are big players providing data that is public and ubiquitous. You may be tempted to push this information off to the side. After all, the vast majority of your company's customer interactions generate private data and not so much at that.

Or do they?

### Launch an App, Create Data

Data generated from customer interaction may have been small a few years ago, but no more. Everyone is seeking new ways to interact with and serve customers, particularly in the mobile realm.

Suppose you work for an insurance company and you want to improve customer service by launching a mobile app that allows users to self-report accidents. For each incident, from the fender bender to the sideswipe to the full on wreck, assume the phone app collects:

- Five high-resolution images of about 1 megabyte (MB) each
- 10 KB worth of meta data for each picture, including the GPS location, time, date and account information
- 10 KB worth of accident descriptions provided by your customer

That adds up to about 5.5 MB for each incident. If 1,000 accidents are reported each day, you need to store about 5 gigabits (GB) of new data every 24 hours. These numbers are still child's play, but it grows when you multiply that number by 365 to give you an annual storage figure of 1.8 Terabytes (TB).

Fewer than 2 TB a year seems manageable. However, that number represents just the customer side of the equation. For each self-submitted claim sent through your new mobile app, your company will probably add a lot more data, including police reports, vehicle registration data and correspondence. Say that adds up to about 20 MB per claim. Now your daily 1,000 mobile-submitted claims generate up to 20 GB of new data. Over the course of a year, that adds up to 7.3 TB.

Together, all the data generated by your single mobile app has now reached 9.1 TB a year. That is data you also have to backup or replicate in some way. Assuming you are a risk-averse organization (you are an insurance company in this example), you plan to keep multiple compressed copies that generate the need for another 4.5 TB of storage annually. Now, your total storage needs to be 13.6 TB.

Data Source	Size (TB)
Customer Data	1.8
Transaction Data	7.3
<b>Subtotal</b>	<b>9.1</b>
Backup	4.55
<b>Total</b>	<b>13.65</b>

Table 1. A simple mobile app can generate more than 13 Tb of data a year.

These kinds of numbers are probably manageable with your existing on-premise SAN, but what happens as the number of customers grows and your app becomes more popular? Your calculus will probably tell you to anticipate steady, moderate growth.

Now suppose a hailstorm sweeps through a half dozen Midwestern cities where you have tens of thousands of customers. After a slight delay due to power outages, usage of your mobile app spikes and your claims triple to 3,000 a day for 10 days. That single incident adds 400 GB of data each day, which increases your data needs by 4 TB over the course of the storm, plus another 2 TB for backup.

Remember, these numbers account for just the data generated by your mobile claims tool. You also have to manage storage for the application itself, which of course is containerized so you can scale based on demand. Count on several more gigabytes of block storage for each instance you spin up, plus back-ups.

### The Ceph Advantage

With just a single app, you can see how your storage needs can quickly grow—and contract—like a big bellow that risks blowing away your budget. If you have a proprietary SAN, you will pay a premium for the storage appliance and each disk you add to your array. If you factored in those costs up front, before you ever launched your new mobile app, you recognize the hefty fixed costs. You also know that if anticipated use of your app exceeds your expectations, expanding your proprietary SAN can run into some big dollars.

With SUSE Enterprise Storage, which is based on the popular open-source Ceph project, you can dramatically lower those fixed costs by using commodity hardware and disks, not proprietary equipment. A typical proprietary disk sold by a SAN vendor can cost about \$1,200 per SSD TB compared with about \$650 for an enterprise-grade commodity SSD. For spinning disks, the savings is nearly 75 percent.

However, it is more than just the cost of the raw hardware that makes Ceph solutions like SUSE Enterprise Storage ideal for providing storage for your mobile app. SUSE Enterprise Storage makes it easy to add new hardware and quickly integrate that new storage into your stack.

New object storage devices (OSDs) are created as you add disks to your cluster and that new storage is immediately available for your workloads. SUSE Enterprise Storage uses Salt and Deepsea, robust automation tools, to deploy and manage your cluster. That makes adding physical storage straightforward. As new servers are

added, new OSDs can be immediately carved up into storage pools and exported as block, object and file stores for your workloads. You can also readily create storage tiers that take advantage of the different kinds of disks you have, from spinners to SSDs.

Replication is built in by your predetermined CRUSH map, which uniformly distributes your data across your cluster. That means there is no single point of failure and no need to set up RAID striping. The cluster automatically mirrors data and rebalances its placement. Later, if your needs change, you can scale down and, again, your data will be uniformly redistributed among the disks remaining in your array.

SUSE® is able to offer the latest features and performance enhancements, support the latest hardware, and do it all while continuing to deliver enterprise reliability and application compatibility. Amazon EC2 customers can run SUSE Linux Enterprise Server on all instance types and sizes, enabling them to maximize cost-effectiveness and performance for a particular workload.

The screenshot shows the SUSE Enterprise Storage dashboard. The top navigation bar includes 'Dashboard', 'OSDs' (highlighted), 'RBDs', 'Pools', 'Nodes', 'iSCSI', 'NFS', 'Object Gateway', 'CRUSH Map', and 'System'. Below the navigation is a 'Ceph OSDs' section with a table of OSDs. The table has columns for Name, Hostname, Status, Storage Backend, and Crush Weight. There are 8 rows of data. Below the table, it says 'Showing 1 to 8 of 8 items'.

Name	Hostname	Status	Storage Backend	Crush Weight
osd.0	data02	down	bluestore	0.038986
osd.1	mon01	up	bluestore	0.038986
osd.2	mon02	up	bluestore	0.038986
osd.3	data01	up	bluestore	0.038986
osd.4	data02	down	bluestore	0.038986
osd.5	data01	up	bluestore	0.038986
osd.6	mon02	up	bluestore	0.038986
osd.7	mon01	up	bluestore	0.038986

**Figure 1.** Object storage devices (OSDs) can be viewed and managed from the dashboard.

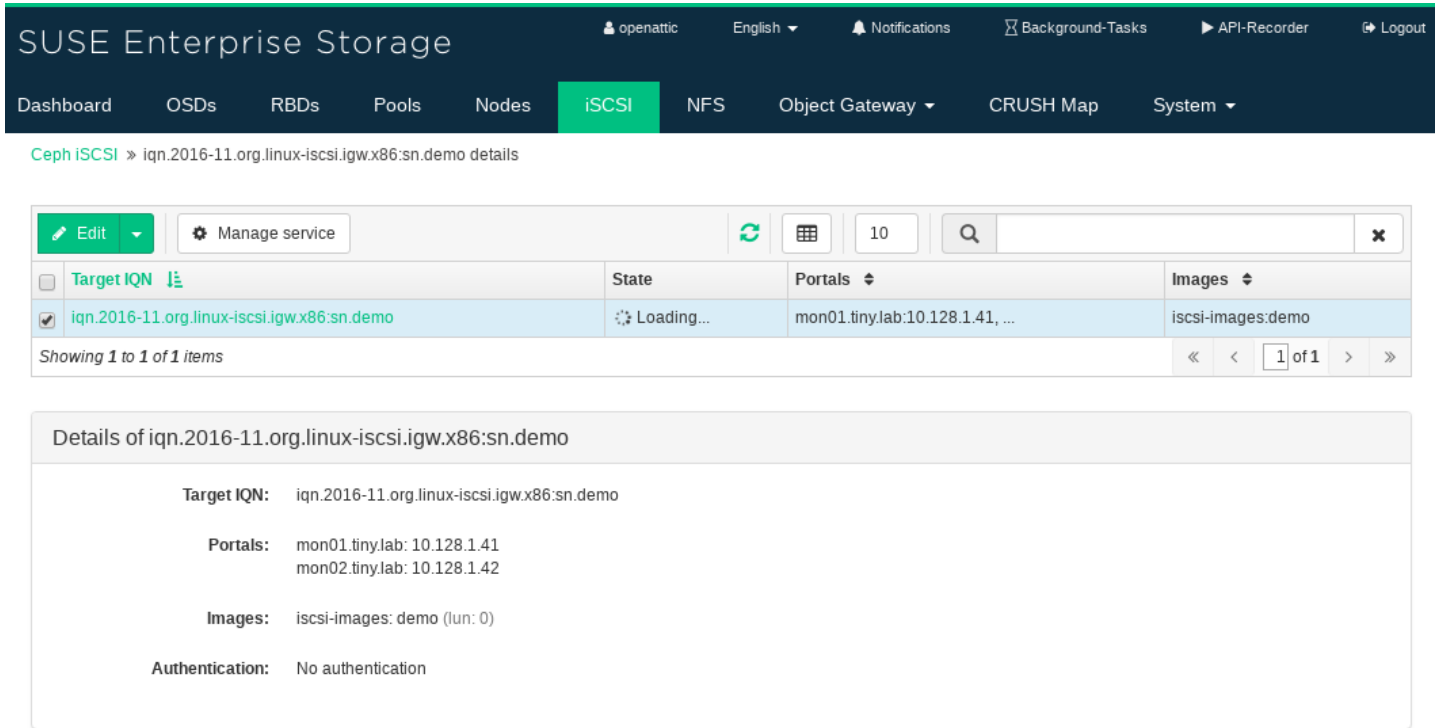


Figure 2. The SUSE Enterprise Storage management console makes it easy to export volumes via iSCSI and other gateways.

### Conclusion

Through this simple scenario, you can begin to develop an idea of how the data explosion can affect your company’s business. As customers continue to demand new and better services, your company will respond with data-driven services and applications

that consume lots of storage. SUSE® is committed to open-source standards and, as a key contributor to Ceph and a founding member of the Ceph Foundation, offers SUSE Enterprise Storage to help you meet those demands, lower your costs and avoid the vendor lock-in that can limit your agility.

Ceph Pools > details

<input type="checkbox"/>	Name	ID	Used	Applications	Placement groups	Replica size	Erasure code profile	Last change	Type	Crush ruleset	Compression mode
<input type="checkbox"/>	.rgw.root	1	0.00%	rgw	8	3		23	replicated	0	none
<input type="checkbox"/>	cephfs_data	6	0.00%	cephfs	128	3		52	replicated	0	none
<input type="checkbox"/>	cephfs_metadata	7	0.00%	cephfs	128	3		52	replicated	0	none
<input type="checkbox"/>	default.rgw.control	2	0.00%	rgw	8	3		28	replicated	0	none
<input type="checkbox"/>	default.rgw.log	4	0.00%	rgw	8	3		35	replicated	0	none
<input type="checkbox"/>	default.rgw.meta	3	0.00%	rgw	8	3		32	replicated	0	none
<input checked="" type="checkbox"/>	iscsi-images	5	0.00%	rbd	128	3		45	replicated	0	none

Showing 1 to 7 of 7 items

Details

Statistics

### Status of iscsi-images

**ID:** 5  
**Pool size:** 98.35 GiB  
**Used:** 35.00 B  
**Applications:** rbd  
**Placement Groups:** 128  
**Replicated size:** 3  
**Type:** replicated  
**Flags:** hashspool  
**Last change:** 45

Figure 3. The SUSE Enterprise Storage management console makes it easy to add pools and RADOS block devices (RBDs).

To learn more, visit  
[suse.com/products/suse-enterprise-storage](https://suse.com/products/suse-enterprise-storage)

[www.suse.com](https://www.suse.com)

