

Open Infrastructure - Task #5786

Document our ceph setup && maintenance guide

10/09/2018 12:14 PM - Nico Schottelius

Status:	Closed	Start date:	10/09/2018
Priority:	Normal	Due date:	
Assignee:	Nico Schottelius	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
PM Check date:			

Description

Introduction

This article describes the ungleich storage architecture that is based on ceph. It describes our architecture as well maintenance commands. Required for

Communication guide

Usually when disks fails no customer communication is necessary, as it is automatically compensated/rebalanced by ceph. However in case multiple disk failures happen at the same time, I/O speed might be reduced and thus customer experience impacted.

For this reason communicate whenever I/O recovery settings are temporarily tuned.

Adding a new disk/ssd

Moving a disk/ssd to another server

Removing a disk/ssd

Handling DOWN osds with filesystem errors

If an email arrives with the subject "monit alert -- Does not exist osd.XX-whoami", the filesystem of an OSD cannot be read anymore. It is very highly likely that the disk / ssd is broken. Steps that need to be done:

- Login to any ceph monitor (cephX.placeY.ungleich.ch)
- Check **ceph -s**, find host using **ceph osd tree**
- Login to the affected host
- Run the following commands:
 - `ls /var/lib/ceph/osd/ceph-XX`
 - `dmesg`
- Create a new ticket in the datacenter light project
 - Subject: "Replace broken OSD.XX on serverX.placeY.ungleich.ch"
 - Add (partial) output of above commands
 - Use `/opt/ungleich-tools/ceph-osd-stop-remove-permanently XX`, where XX is the osd id, to remove the disk from the cluster
 - Remove the physical disk from the host, checkout if there is warranty on it and if yes
 - Create a short letter to the vendor, including technical details a from above
 - Record when you sent it in
 - Put ticket into status waiting
 - If there is no warranty, dispose it

Change ceph speed for i/o recovery

By default we want to keep I/O recovery traffic low to not impact customer experience. However when multiple disks fail at the same point, we might want to prioritise recover for data safety over performance.

The default configuration on our servers contains:

```
[osd]
osd max backfills = 1
```

```
osd recovery max active = 1
osd recovery op priority = 2
```

The important settings are **osd max backfills** and **osd recovery max active**, the priority is always kept low so that regular I/O has priority.

To adjust the number of backfills **per osd** and to change the **number of threads** used for recovery, we can use on any node with the admin keyring:

```
ceph tell osd.* injectargs '--osd-max-backfills Y'
ceph tell osd.* injectargs '--osd-recovery-max-active X'
```

where Y and X are the values that we want to use. Experience shows that Y=5 and X=5 doubles to triples the recovery performance, whereas X=10 and Y=10 increases recovery performance 5 times.

Debug scrub errors / inconsistent pg message

From time to time disks don't save what they are told to save. Ceph scrubbing detects these errors and switches to HEALTH_ERR. Use **ceph health detail** to find out which placement groups (**pgs**) are affected. Usually a `*ceph pg repair <number>` fixes the problem.

If this does not help, consult <https://ceph.com/een-categorie/ceph-manually-repair-object/>.

History

#1 - 10/09/2018 12:20 PM - Nico Schottelius

- Description updated

- Status changed from New to In Progress

Will put sample / first versions in this ticket before migrating to the wiki.

#2 - 10/09/2018 12:24 PM - Nico Schottelius

- Description updated

#3 - 10/09/2018 12:26 PM - Nico Schottelius

- Description updated

#4 - 10/09/2018 12:37 PM - Nico Schottelius

- Description updated

#5 - 10/09/2018 12:42 PM - Nico Schottelius

- Description updated

#6 - 10/09/2018 01:07 PM - Nico Schottelius

- Description updated

#7 - 10/19/2018 01:00 PM - Samuel Hailu

Nico Schottelius wrote:

Introduction

This article describes the ungleich storage architecture that is based on ceph. It describes our architecture as well maintenance commands. Required for

Communication guide

Usually when disks fails no customer communication is necessary, as it is automatically compensated/rebalanced by ceph. However in case multiple disk failures happen at the same time, I/O speed might be reduced and thus customer experience impacted.

For this reason communicate whenever I/O recovery settings are temporarily tuned.

Note please inform or notify in the infrastructure channel before you add or remove a disk

h2. Adding a new disk/ssd

Moving a disk/ssd to another server

Removing a disk/ssd

Change ceph speed for i/o recovery

By default we want to keep I/O recovery traffic low to not impact customer experience. However when multiple disks fail at the same point, we might want to prioritise recover for data safety over performance.

The default configuration on our servers contains:

[...]

The important settings are **osd max backfills** and **osd recovery max active**, the priority is always kept low so that regular I/O has priority.

To adjust the number of backfills **per osd** and to change the **number of threads** used for recovery, we can use on any node with the admin keyring:

[...]

where Y and X are the values that we want to use. Experience shows that Y=5 and X=5 doubles to triples the recovery performance, whereas X=10 and Y=10 increases recovery performance 5 times.

Debug scrub errors / inconsistent pg message

From time to time disks don't save what they are told to save. Ceph scrubbing detects these errors and switches to HEALTH_ERR. Use **ceph health detail** to find out which placement groups (**pgs**) are affected. Usually a `*ceph pg repair <number>` fixes the problem.

If this does not help, consult <https://ceph.com/geen-categorie/ceph-manually-repair-object/>.

Nico Schottelius wrote:

Introduction

This article describes the ungleich storage architecture that is based on ceph. It describes our architecture as well maintenance commands. Required for

Communication guide

Usually when disks fails no customer communication is necessary, as it is automatically compensated/rebalanced by ceph. However in case multiple disk failures happen at the same time, I/O speed might be reduced and thus customer experience impacted.

For this reason communicate whenever I/O recovery settings are temporarily tuned.

Adding a new disk/ssd

Moving a disk/ssd to another server

Removing a disk/ssd

Change ceph speed for i/o recovery

By default we want to keep I/O recovery traffic low to not impact customer experience. However when multiple disks fail at the same point, we might want to prioritise recover for data safety over performance.

The default configuration on our servers contains:

[...]

The important settings are **osd max backfills** and **osd recovery max active**, the priority is always kept low so that regular I/O has priority.

To adjust the number of backfills **per osd** and to change the **number of threads** used for recovery, we can use on any node with the admin keyring:

[...]

where Y and X are the values that we want to use. Experience shows that Y=5 and X=5 doubles to triples the recovery performance, whereas X=10 and Y=10 increases recovery performance 5 times.

Debug scrub errors / inconsistent pg message

From time to time disks don't save what they are told to save. Ceph scrubbing detects these errors and switches to HEALTH_ERR. Use **ceph health detail** to find out which placement groups (pgs) are affected. Usually a `*ceph pg repair <number>` fixes the problem.

If this does not help, consult <https://ceph.com/geen-categorie/ceph-manually-repair-object/>.

#8 - 10/19/2018 01:25 PM - Samuel Hailu

Nico Schottelius wrote:

Introduction

This article describes the ungleich storage architecture that is based on ceph. It describes our architecture as well maintenance commands. Required for

Communication guide

Usually when disks fails no customer communication is necessary, as it is automatically compensated/rebalanced by ceph. However in case multiple disk failures happen at the same time, I/O speed might be reduced and thus customer experience impacted.

For this reason communicate whenever I/O recovery settings are temporarily tuned.

Note please inform or notify in the infrastructure channel before you add or remove a disk.

Adding a new disk/ssd

Adding a new disk/ssd

Moving a disk/ssd to another server

Removing a disk/ssd

Change ceph speed for i/o recovery

By default we want to keep I/O recovery traffic low to not impact customer experience. However when multiple disks fail at the same point, we might want to prioritise recover for data safety over performance.

The default configuration on our servers contains:

[...]

The important settings are **osd max backfills** and **osd recovery max active**, the priority is always kept low so that regular I/O has priority.

To adjust the number of backfills **per osd** and to change the **number of threads** used for recovery, we can use on any node with the admin keyring:

[...]

where Y and X are the values that we want to use. Experience shows that Y=5 and X=5 doubles to triples the recovery performance, whereas X=10 and Y=10 increases recovery performance 5 times.

Debug scrub errors / inconsistent pg message

From time to time disks don't save what they are told to save. Ceph scrubbing detects these errors and switches to HEALTH_ERR. Use **ceph health detail** to find out which placement groups (pgs) are affected. Usually a `*ceph pg repair <number>` fixes the problem.

If this does not help, consult <https://ceph.com/geen-categorie/ceph-manually-repair-object/>.

#9 - 10/19/2018 01:27 PM - Samuel Hailu

Nico Schottelius wrote:

Introduction

This article describes the ungleich storage architecture that is based on ceph. It describes our architecture as well maintenance commands.

Required for

Communication guide

Usually when disks fails no customer communication is necessary, as it is automatically compensated/rebalanced by ceph. However in case multiple disk failures happen at the same time, I/O speed might be reduced and thus customer experience impacted.

For this reason communicate whenever I/O recovery settings are temporarily tuned.

Note please inform or notify in the infrastructure channel **before** you add or remove a disk.

Adding a new disk/ssd

Adding a new disk/ssd

Moving a disk/ssd to another server

Removing a disk/ssd

Change ceph speed for i/o recovery

By default we want to keep I/O recovery traffic low to not impact customer experience. However when multiple disks fail at the same point, we might want to prioritise recover for data safety over performance.

The default configuration on our servers contains:

[...]

The important settings are **osd max backfills** and **osd recovery max active**, the priority is always kept low so that regular I/O has priority.

To adjust the number of backfills **per osd** and to change the **number of threads** used for recovery, we can use on any node with the admin keyring:

[...]

where Y and X are the values that we want to use. Experience shows that Y=5 and X=5 doubles to triples the recovery performance, whereas X=10 and Y=10 increases recovery performance 5 times.

Debug scrub errors / inconsistent pg message

From time to time disks don't save what they are told to save. Ceph scrubbing detects these errors and switches to HEALTH_ERR. Use **ceph health detail** to find out which placement groups (**pgs**) are affected. Usually a `*ceph pg repair <number>` fixes the problem.

If this does not help, consult <https://ceph.com/geen-categorie/ceph-manually-repair-object/>.

#10 - 10/20/2018 02:11 PM - Nico Schottelius

- *Description updated*

#11 - 12/29/2018 12:55 PM - Nico Schottelius

- *Status changed from In Progress to Closed*