

TECH WEEK STORAGE 24



EOS Open Storage

III - Instance Inventory Implementation

Dr. Andreas-Joachim Peters

for the EOS Project - CERN IT - Storage Group

IT Auditorium - CERN

15.03.2024

- Why do we need an instance inventory?
- Device Registry
- Instance Value
- Data Storage Cost
- Summary & Outlook



Why III ?

Instance Inventory Implementation



Why III ?

Instance Inventory Implementation

- **EOS** instances have until today already a hllifetime of a decade





Why III ?

Instance Inventory Implementation

- **EOS** instances have until today already a hllifetime of a decade
- over time a lot of hardware is replaced - **historic information is lost**





Why III ?

Instance Inventory Implementation

- **EOS** instances have until today already a hllifetime of a decade
 - over time a lot of hardware is replaced - **historic information is lost**
 - **EOS** itself did not provide an overview about the **age/model** of storage devices installed





Why III ?

Instance Inventory Implementation

- **EOS** instances have until today already a lifetime of a decade
 - over time a lot of hardware is replaced - **historic information is lost**
- **EOS** itself did not provide an overview about the **age/model** of storage devices installed
 - we only provided number of disks and volume in EOS views





Why III ?

Instance Inventory Implementation

- **EOS** instances have until today already a lifetime of a decade
 - over time a lot of hardware is replaced - **historic information is lost**
- **EOS** itself did not provide an overview about the **age/model** of storage devices installed
 - we only provided number of disks and volume in EOS views





Why III ?

Instance Inventory Implementation



- **EOS** instances have until today already a hllifetime of a decade
 - over time a lot of hardware is replaced - **historic information is lost**
 - **EOS** itself did not provide an overview about the **age/model** of storage devices installed
 - we only provided number of disks and volume in EOS views
- an instance wide **accounting** based on **duration** how long files have already been stored was missing
 - File/Volume usage is available in the Quota accounting



Why III ?

Instance Inventory Implementation



- **EOS** instances have until today already a lifetime of a decade
 - over time a lot of hardware is replaced - **historic information is lost**
 - **EOS** itself did not provide an overview about the **age/model** of storage devices installed
 - we only provided number of disks and volume in EOS views
- an instance wide **accounting** based on **duration** how long files have already been stored was missing
 - File/Volume usage is available in the Quota accounting
- **EOS** didn't provide a **good estimator for cost & value** of data and



Device Registry



- Each device which has ever had an assigned filesystem id in **EOS** is now forever tracked under

`/eos/<instance>/proc/devices/<device-id>.<fs-id>`

```
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16NF95G.27779
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16SF95G.27786
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A19YF95G.27768
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1CXF95G.27778
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FMF95G.27784
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FNF95G.27788
rw-r--r-- 0 root root 0 Nov 28 13:20 6970A01GF95G.27776
```

Device Registry

- Each device which has ever had an assigned filesystem id in **EOS** is now forever tracked under

`/eos/<instance>/proc/devices/<device-id>.<fs-id>`

```
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16NF95G.27779
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16SF95G.27786
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A19YF95G.27768
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1CXF95G.27778
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FMF95G.27784
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FNF95G.27788
rw-r--r-- 0 root root 0 Nov 28 13:20 6970A01GF95G.27776
```

- **btime** when device/id pair was first seen

- Each device which has ever had an assigned filesystem id in **EOS** is now forever tracked under

`/eos/<instance>/proc/devices/<device-id>.<fs-id>`

```
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16NF95G.27779
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16SF95G.27786
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A19YF95G.27768
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1CXF95G.27778
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FMF95G.27784
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FNF95G.27788
rw-r--r-- 0 root root 0 Nov 28 13:20 6970A01GF95G.27776
```

- **btime** when device/id pair was first seen
- **mtime** last time a complete JSON smart status was stored

- Each device which has ever had an assigned filesystem id in **EOS** is now forever tracked under

`/eos/<instance>/proc/devices/<device-id>.<fs-id>`

```
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16NF95G.27779
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16SF95G.27786
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A19YF95G.27768
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1CXF95G.27778
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FMF95G.27784
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FNF95G.27788
rw-r--r-- 0 root root 0 Nov 28 13:20 6970A01GF95G.27776
```

- **btime** when device/id pair was first seen
- **mtime** last time a complete JSON smart status was stored
- extended attribute **sys.smart.status**

- Each device which has ever had an assigned filesystem id in **EOS** is now forever tracked under

`/eos/<instance>/proc/devices/<device-id>.<fs-id>`

```
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16NF95G.27779
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16SF95G.27786
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A19YF95G.27768
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1CXF95G.27778
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FMF95G.27784
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FNF95G.27788
rw-r--r-- 0 root root 0 Nov 28 13:20 6970A01GF95G.27776
```

- **btime** when device/id pair was first seen
- **mtime** last time a complete JSON smart status was stored
- extended attribute **sys.smart.status**
 - **OK, NOCTL, NA, FAILING, CHECK, INVAL, UNKNOWN**

- Each device which has ever had an assigned filesystem id in **EOS** is now forever tracked under

`/eos/<instance>/proc/devices/<device-id>.<fs-id>`

```
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16NF95G.27779
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16SF95G.27786
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A19YF95G.27768
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1CXF95G.27778
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FMF95G.27784
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FNF95G.27788
rw-r--r-- 0 root root 0 Nov 28 13:20 6970A01GF95G.27776
```

- **btime** when device/id pair was first seen
- **mtime** last time a complete JSON smart status was stored
- extended attribute **sys.smart.status**
 - **OK, NOCTL, NA, FAILING, CHECK, INVAL, UNKNOWN**
- extended attribute **sys.smart.json**

- Each device which has ever had an assigned filesystem id in **EOS** is now forever tracked under

`/eos/<instance>/proc/devices/<device-id>.<fs-id>`

```
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16NF95G.27779
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16SF95G.27786
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A19YF95G.27768
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1CXF95G.27778
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FMF95G.27784
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FNF95G.27788
rw-r--r-- 0 root root 0 Nov 28 13:20 6970A01GF95G.27776
```

- **btime** when device/id pair was first seen
- **mtime** last time a complete JSON smart status was stored
- extended attribute **sys.smart.status**
 - **OK, NOCTL, NA, FAILING, CHECK, INVAL, UNKNOWN**
- extended attribute **sys.smart.json**
 - output of `smartctl -x -j <device>`

- Each device which has ever had an assigned filesystem id in **EOS** is now forever tracked under


`/eos/<instance>/proc/devices/<device-id>.<fs-id>`

```
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16NF95G.27779
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A16SF95G.27786
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A19YF95G.27768
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1CXF95G.27778
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FMF95G.27784
rw-r--r-- 0 root root 0 Nov 28 13:20 59F0A1FNF95G.27788
rw-r--r-- 0 root root 0 Nov 28 13:20 6970A01GF95G.27776
```

- **btime** when device/id pair was first seen
- **mtime** last time a complete JSON smart status was stored
- extended attribute **sys.smart.status**
 - **OK, NOCTL, NA, FAILING, CHECK, INVAL, UNKNOWN**
- extended attribute **sys.smart.json**
 - output of `smartctl -x -j <device>`

```
EOS Console [root://localhost] |/eos/cms/proc/devices/> attr get sys.smart.status Z6L4K81EFTMB.23300
sys.smart.status="OK"
```

Device S.M.A.R.T. information collection



FST

FST

Shared Hash

MGM




```
smartctl -x -j
```

```
smartctl -x -j
```

```
smartctl -x -j
```

FST



```
smartctl -x -j
```

FST

Shared Hash


MGM



smartctl -x -j

smartctl -x -j

smartctl -x -j



smartctl -x -j

```
stat.health.z64smart := zbase64:MDAwMGEzMjllN01dSXp0B0Ba+51egd0qshyuktg3t6zEqfGSKZROzXS lWDAJSexwKy6W3V357w0AG0iBWixaNjvMTaYTEHJA9zaAD49/vw0g92fouf rCCxwY6i C3P7f0K/sbFAIhn91+A//96LnuhA4PI1Gxc429a1Ks+AsDvLL006N9vyc9e+0DqAXq0to92TSSezmZ80IkkBb7sIjhy+ jgT5Q+fb Lxo99+VwUzom+KA6E8+HgX0xH9vA8qZE1cB9btpk//YvtGdAG90a/s1owMD4wVB jT9R+Z6z+L6/cmengfmbXKw01PVqRHkYw1KpCLhnsDzPhuLzLOGGK0100wza00EH2JfPwxu5h/yypETz4iDmGJT1bgRZ7HEQ6F/OLXt645n1zhv8ejk640uXq6EKqjct0ZXCFT1QTRVEFCwXbuULCqatdXV4P mxeu1yKTP8067dqLlLZANq0T/zZfQZAFAR1baK7Lwqun0w0vAK986uP11+LayHb3z2AGAb0mFFFT0m9x1772R0p0sK90B0wJ D05nCJESURREQRW1sKL81Dv0v0mFrZdy3ooUXm02V1G1h3oKw8v0RvQ MWJk6gMSMF0eARhIQ+FCtVKCqEUIZ0f4B2T0pKZ2WU5XMYWk6xwp1a1faunnGbbuKH7ZFT616pHSJagTayT1TAsXqFLrukK19nFekMYLBFXWANS4L8Xkxw08SZ8TFc1Fw4J9ekJmCpK0FVguUGCIX0hjj u4xP14A+ s1HrqMECMU4HLceugjZBN0PAX/5H3lRn50Jw9D4F8DH+/dhIX+xa0Wh7mN2CZHuaTxgZCX3meFpBodcfqPVv1R0mzEQYdcqYU+eUpYR1VMGJk0Yy5IL90+1BQEMWhMxLsMH0SKv25jECT94DKEZALH+V5V8FAU
```




compressed + b64 encoded JSON
at startup of FST and then every
hour



MGM



smartctl -x -j
smartctl -x -j
smartctl -x -j



smartctl -x -j

```
stat.health.z64smart := zbase64:MDAwMGEzMj14n01dSXp0Ba+51egd0qshyuktg3t6zEqfGSKZROzXS1WDAJSexwKy6W3V357w0AG01BWixaNjvMTaYTEHJA9zaAD49/vw0g92fouf rCCxwY6i  
C3P7f0K/sbFAIhn91+A//96LnuhA4PI16xc429a1Ks+AsDvLL006N9vyc9e+0DqAXq0to92TSSezmZb80IkkBb7s1JhY+ jgTSQ+fb Lxo99+VwUzom+KA6E8+HgX0xH9vA8qZE1cB9btpk//YvtGdAG9oa/siowM4wVB  
jt9R+Z6z+L6/cmengfmbxKw01PVqRHkYw1kPCLhnsDzPhuLZ1oGKGX01000za00EH2JfWxu5h/yypETz4tDmGJT1bgRZ7hEQ6F/OLXt645n1zhv8eJk40uXq6ERkQjyct0ZXCFT1Q1RVFfcmXbuULCqatdXV4P  
mxeu1yXZPB0617dqlLlFal1zZn0qT/zZfQZAFmR1baK7Lwqun0w0vAK96buP11+1ayHb3z2k6Ab0mFFt0m9x1772R0p0s0000wJ005nCJESURREQRW1sKL81DvVv0vFrZdy3oo0UXm02V1G1h1z0Kw8v0RvQ  
MMJkEGMSMF0eARhIQ+FCtVKCqE1Z0f4B2T0pKZ2WU5XMYK6xwp1a1faunnGbbuKH72F1616ghSJa9TayT1TAsXqFlrukK19nFekMYLBFxWNSL8XhoW8SZ8FfC1Fw4J9ekJmCpK0FVguUGC1OXn1j u4xP14A+  
siHrqMECMU4HLceugjZBN0PAX/5H31Rn50Jw9D4F8DH+/dhIX+xa0Wh7mN2CZHuaTxgZCX3meFpBodcfqPvV1R0mzEQYdcqJYU+eUpYR1VMGJk0Yy5IL90+1BQEMWhMxLsMH0SKv25jECT94DkEZA1H+V5V8FAU
```



compressed + b64 encoded JSON
at startup of FST and then every
hour



MGM


Hard disk icons created by vectorsmarket15 - Flaticon



smartctl -x -j

smartctl -x -j

smartctl -x -j

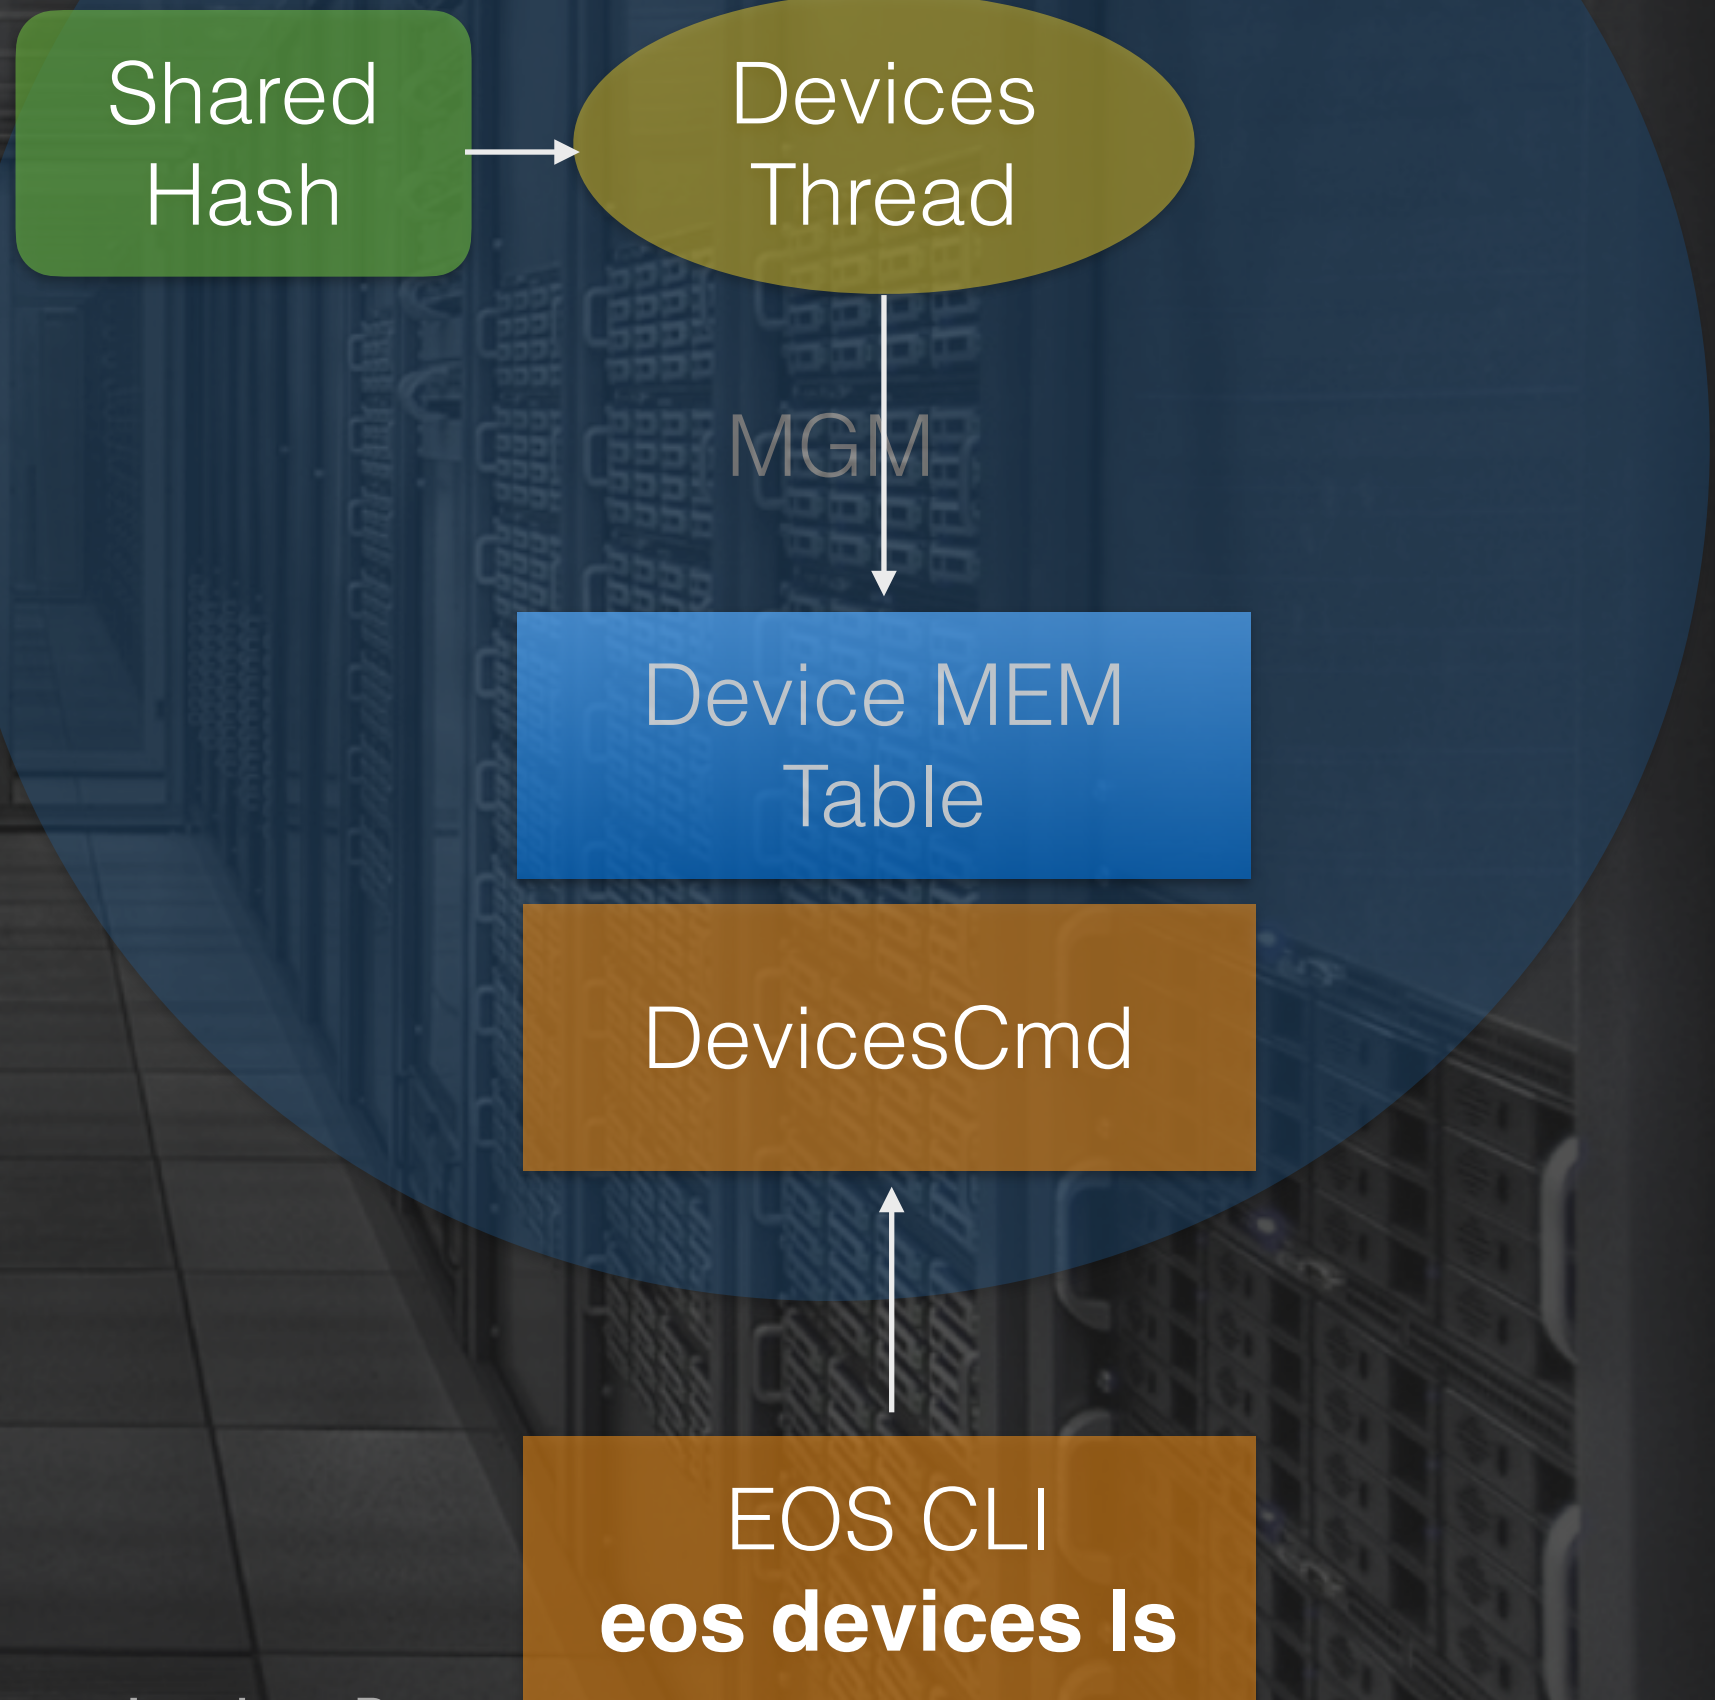


smartctl -x -j

```
stat.health.z64smart := zbase64:MDAwMGEzMj14n01dSXp0Ba+51eqd0qshyuktg3t6zEqfGSKZROzXS1WDAJSexwKy6W3V357w0AG01BWixaNjvMTaYTEHJA9zaAD49/vw0g92fourfCCxwY6iC3P7f0K/sbFAIhn91+A//96LnuhA4PI16xc429a1Ks+AsDvLL006N9vyc9e+0DqAXq0to9ZTSSezmZb80IkKbB7s1JhY+ jgTSQ+fb Lxo99+VwUzom+KA6E8+HgX0xH9vA8qZE1cB9btpk//YvtGdAG9oa/siowMD4wBjT9R+Z6z+L6/cmengfmbXKw01PVqRHkYw1kPCLhnsDzPhuLZLOGKX010w0za00EH2JFwXu5h/yypETz4tDmGUT1bgRZ7hEQ6F/OLXt645n1zhv8eJk40uQ6E8kQjct0ZXCft1Q1RVFfCWbuULCqatdXV4Pmexu1yXZPB0617dqlLlFalLzZk0qT/zfzQZAFmR1baK7Lwqun0w0vAK966uP11+*ayHb3z2k6Ab0mFFt0m9x1772R0p0s0000w1D05nCJESURREQRW1sKL81Dv0v0mFRZdy3ooUXm02V1G1h3oKw8v0RvQWJ3k6GMSMF0eARhIQ+FC1VKCqE1Z0F4B2T0pKZ2WU5XMYK6xwp1a1faunnGbuKH72F1616pHSJagTayT1TAsXqFLrukK19nFekMYLBFxAN5qL8XhoW8SZ8FfC1Fw4J9ekJmCpK0FV0UGC1OXh1j4xP14A+siHrqMECMU4HLceugjZBN0PAX/5H3lRn50Jw9D4F8DH+/dhIX+xa0Wh7m2CZHuaTxgZCX3meFpBodcfqPVv1R0mzEQYdcqYU+eUpYR1VMGJk0Yy5IL90+1BQEMWhMxLsMH0SKv25jECT94DkEZA1H+V5V8FAU
```




compressed + b64 encoded JSON
at startup of FST and then every hour





EOS Devices CLI




Summary mode: **eos devices ls**

Summary mode: **eos devices ls**

space	model	avg:age [years]	bytes	count	hours	smrt:ok	smrt:noctl	smrt:na	smrt:failing	smrt:check	smrt:inval	smrt:unknown
default	HGST:HMS5C4040BLE640	8.70	900.18 TB	225	17.14 Mh	225	0	0	0	0	0	0
default	HGST:HUH721212ALE600	4.50	6.23 PB	519	20.46 Mh	519	0	0	0	0	0	0
default	HGST:HUS724020ALA640	8.74	26.01 TB	13	995.09 Kh	13	0	0	0	0	0	0
default	HGST:HUS726040ALE610	2.32	4.00 TB	1	20.33 Kh	1	0	0	0	0	0	0
default	HGST:HUS726060ALE610	6.51	10.53 PB	1.75 K	100.07 Mh	1754	0	0	0	0	0	0
default	HGST:HUS726060ALE614	4.41	42.01 TB	7	270.25 Kh	7	0	0	0	0	0	0
default	HGST:HUS726T6TALE6L4	3.54	24.00 TB	4	124.03 Kh	4	0	0	0	0	0	0
default	INTEL:SSDSC2BB960G7	6.60	48.97 TB	51	2.95 Mh	51	0	0	0	0	0	0
default	ST4000NC001-1FS168	8.03	640.13 TB	160	11.25 Mh	160	0	0	0	0	0	0
default	ST4000NM0035-1V4107	5.30	20.00 TB	5	232.36 Kh	5	0	0	0	0	0	0
default	ST4000NM0115-1YZ107	5.28	4.00 TB	1	46.23 Kh	1	0	0	0	0	0	0
default	TOSHIBA:MG04ACA600E	6.53	10.08 PB	1.68 K	96.06 Mh	1679	0	0	0	0	0	0
default	TOSHIBA:MG06ACA600E	3.64	12.00 TB	2	63.77 Kh	2	0	0	0	0	0	0
default	TOSHIBA:MG07ACA12TE	3.90	47.28 PB	3.94 K	134.56 Mh	3940	0	0	0	0	0	0
default	TOSHIBA:MG07ACA14TE	2.25	4.00 PB	286	5.62 Mh	286	0	0	0	0	0	0
default	WDC::WUH721414ALE6L4	2.11	25.31 PB	1.81 K	33.46 Mh	1808	0	0	0	0	0	0
default	WDC::WUH721818ALE6L4	0.91	3.44 PB	191	1.51 Mh	191	0	0	0	0	0	0



EOS Devices CLI



Details mode: **eos devices ls -l**

Details mode: **eos devices ls -l**

default	model	serial	type	capacity	rpms	poweron[h]	temp[degrees]	S.M.A.R.T	if	rla	wc
	INTEL:SSDSC2BB960G7	PHDV648501Z6960FGN	sat	960.20 GB	0	58.06 Kh	18	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	NCHEV6BS	sat	6.00 TB	7200	57.75 Kh	29	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1HAP7DB	sat	6.00 TB	7200	57.75 Kh	28	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1HD72WD	sat	6.00 TB	7200	57.75 Kh	29	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	NCHE4SUS	sat	6.00 TB	7200	57.75 Kh	29	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	NCHEV7ZS	sat	6.00 TB	7200	57.75 Kh	30	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1HEEHVD	sat	6.00 TB	7200	57.75 Kh	30	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1HE412D	sat	6.00 TB	7200	57.75 Kh	31	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1HBS9ZB	sat	6.00 TB	7200	57.75 Kh	31	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	NCHEV4HS	sat	6.00 TB	7200	57.75 Kh	30	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1HDMBED	sat	6.00 TB	7200	57.75 Kh	31	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	NCHEV61S	sat	6.00 TB	7200	57.75 Kh	30	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	NCHENHSS	sat	6.00 TB	7200	57.75 Kh	31	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1H9D21D	sat	6.00 TB	7200	57.75 Kh	29	ok	6.0:Gb/s	true	true
	INTEL:SSDSC2BB960G7	PHDV648502JC960FGN	sat	960.20 GB	0	58.06 Kh	19	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1HAN32B	sat	6.00 TB	7200	57.75 Kh	30	ok	6.0:Gb/s	true	true
	HGST:HUS726060ALE610	K1HE8Y0D	sat	6.00 TB	7200	57.75 Kh	30	ok	6.0:Gb/s	true	true

Instance **Value** output: **eos devices ls**

Cost-Matrix	TB*Years	Avg-Drive-Hours	Tot-Drive-Hours	Cloud\$-Replica	Cloud\$-Erasure
	426.26 K	39.91 K	424.84 M	53.28 M\$	88.81 M\$

Instance **Value** output: **eos devices ls**

Cost-Matrix	TB*Years	Avg-Drive-Hours	Tot-Drive-Hours	Cloud\$-Replica	Cloud\$-Erasure
	426.26 K	39.91 K	424.84 M	53.28 M\$	88.81 M\$

- computes the **instance value** by using drive hours, capacity and cloud storage prices when used with replication or erasure coding layouts

Instance **Value** output: **eos devices ls**

Cost-Matrix	TB*Years	Avg-Drive-Hours	Tot-Drive-Hours	Cloud\$-Replica	Cloud\$-Erasure
	426.26 K	39.91 K	424.84 M	53.28 M\$	88.81 M\$

- computes the **instance value** by using drive hours, capacity and cloud storage prices when used with replication or erasure coding layouts
 - this is equivalent to the price paid to a cloud provider for the provided space while drives were on

Instance **Value** output: **eos devices ls**

Cost-Matrix	TB*Years	Avg-Drive-Hours	Tot-Drive-Hours	Cloud\$-Replica	Cloud\$-Erasure
	426.26 K	39.91 K	424.84 M	53.28 M\$	88.81 M\$

- computes the **instance value** by using drive hours, capacity and cloud storage prices when used with replication or erasure coding layouts
 - this is equivalent to the price paid to a cloud provider for the provided space while drives were on
- the longer disks are running, the more value they create!

Instance **Value** output: **eos devices ls**

Cost-Matrix	TB*Years	Avg-Drive-Hours	Tot-Drive-Hours	Cloud\$-Replica	Cloud\$-Erasure
	426.26 K	39.91 K	424.84 M	53.28 M\$	88.81 M\$

- computes the **instance value** by using drive hours, capacity and cloud storage prices when used with replication or erasure coding layouts
 - this is equivalent to the price paid to a cloud provider for the provided space while drives were on
- the longer disks are running, the more value they create!
 - cloud price in \$/TB/year configurable (default 250 = approx. AZURE/GOOGLE price in 100PB range)

Instance **Value** output: **eos devices ls**

Cost-Matrix	TB*Years	Avg-Drive-Hours	Tot-Drive-Hours	Cloud\$-Replica	Cloud\$-Erasure
	426.26 K	39.91 K	424.84 M	53.28 M\$	88.81 M\$

- computes the **instance value** by using drive hours, capacity and cloud storage prices when used with replication or erasure coding layouts
 - this is equivalent to the price paid to a cloud provider for the provided space while drives were on
- the longer disks are running, the more value they create!
 - cloud price in \$/TB/year configurable (default 250 = approx. AZURE/GOOGLE price in 100PB range)
 - replication factor configurable (default 2)

Instance **Value** output: **eos devices ls**

Cost-Matrix	TB*Years	Avg-Drive-Hours	Tot-Drive-Hours	Cloud\$-Replica	Cloud\$-Erasure
	426.26 K	39.91 K	424.84 M	53.28 M\$	88.81 M\$

- computes the **instance value** by using drive hours, capacity and cloud storage prices when used with replication or erasure coding layouts
 - this is equivalent to the price paid to a cloud provider for the provided space while drives were on
- the longer disks are running, the more value they create!
 - cloud price in \$/TB/year configurable (default 250 = approx. AZURE/GOOGLE price in 100PB range)
 - replication factor configurable (default 2)
 - erasure coding overhead configurable (default 1.2 = RS 10+2)

[10th of march]

Instance	Drive Hours	TB*Years
eosams	92 Mh	132k
eosalice	440 Mh	441k
eosaliceo2	222 Mh	364k
eosatlas	416 Mh	388k
eoscms	344 Mh	326k
eoslhcb	166 Mh	207k
eospilot	18 Mh	17k
eospublic	365 Mh	357k
SUM	2,061 Gh	2,232M



EOS Data Storage Inventory



How to bill users ... using `eos inspector ls`

How to bill users ... using `eos inspector ls`

- **file inspector** component uses **namespace explorer** to create full instance statistics [explorer interval is configurable] ... scans meta-data of all files

How to bill users ... using `eos inspector ls`

- **file inspector** component uses **namespace explorer** to create full instance statistics [explorer interval is configurable] ... scans meta-data of all files
 - **sum of bytes** stored in each layout

How to bill users ... using `eos inspector ls`

- **file inspector** component uses **namespace explorer** to create full instance statistics [explorer interval is configurable] ... scans meta-data of all files
 - **sum of bytes** stored in each layout
 - **access time distributions** for #files and volume

How to bill users ... using `eos inspector ls`

- **file inspector** component uses **namespace explorer** to create full instance statistics [explorer interval is configurable] ... scans meta-data of all files
 - **sum of bytes** stored in each layout
 - **access time distributions** for #files and volume
 - **birth time distributions** for #files and volume

How to bill users ... using `eos inspector ls`

- **file inspector** component uses **namespace explorer** to create full instance statistics [explorer interval is configurable] ... scans meta-data of all files
 - **sum of bytes** stored in each layout
 - **access time distributions** for #files and volume
 - **birth time distributions** for #files and volume
 - **birth vs access time distributions** for #files and volume

How to bill users ... using `eos inspector ls`

- **file inspector** component uses **namespace explorer** to create full instance statistics [explorer interval is configurable] ... scans meta-data of all files
 - **sum of bytes** stored in each layout
 - **access time distributions** for #files and volume
 - **birth time distributions** for #files and volume
 - **birth vs access time distributions** for #files and volume
 - **cost** of currently stored data by
 - **user**
 - *on disk*
 - *on tape*
 - **group**
 - *on disk*
 - *on tape*

How to bill users ... using `eos inspector ls`

- **file inspector** component uses **namespace explorer** to create full instance statistics [explorer interval is configurable] ... scans meta-data of all files
 - **sum of bytes** stored in each layout
 - **access time distributions** for #files and volume
 - **birth time distributions** for #files and volume
 - **birth vs access time distributions** for #files and volume
 - **cost** of currently stored data by
 - **user**
 - *on disk*
 - *on tape*
 - **group**
 - *on disk*
 - *on tape*

$$\mathbf{COST}/\text{file} = \mathbf{Age}(y) * \mathbf{PhysicalSize}[\text{TB}] * \mathbf{Price} [\$/\text{TB}/y]$$

Volume output: eos inspector ls

```
=====
```

```
Storage Bytes - User View [ disk ]
```

```
-----
```

```
Total Bytes : 80.63 PB
```

```
-----
```

```
01. atlas003           : 53.63 PB
```

```
02. atlascdr          : 10.30 PB
```

```
03. atlas001         : 9.41 PB
```

```
04. atlas001         : 613.64 TB
```

```
05. atlas001         : 320.16 TB
```

```
06. atlas001         : 301.88 TB
```

```
07. atlas001         : 199.13 TB
```

```
08. atlas001         : 147.88 TB
```

```
09. atlas001         : 133.71 TB
```

```
10. beastsrv         : 120.85 TB
```

Cost output: eos inspector ls

```
=====
Storage Costs - User View [ disk ]
-----
Total Costs : 195.02 kEOS
-----
01. atlas003           : 112.10 kEOS
02. wguan              : 17.70 kEOS
03. atlas001          : 6737 EOS
04. jt                 : 4213 EOS
05. jb                 : 1736 EOS
06. 40                 : 1683 EOS
07. 77                 : 1519 EOS
08. 16                 : 1213 EOS
09. 16                 : 1177 EOS
10. cluster            : 1137 EOS
```

Cost output: eos inspector ls

```
=====
Storage Costs - User View [ disk ]
-----
Total Costs : 195.02 kEOS
-----
01. atlas003           : 112.10 kEOS
02. wguan              : 17.70 kEOS
03. atlas001          : 6737 EOS
04. jt                 : 4213 EOS
05. jb                 : 1736 EOS
06. 40                 : 1683 EOS
07. 77                 : 1519 EOS
08. 16                 : 1213 EOS
09. 16                 : 1177 EOS
10. clscr              : 1137 EOS
=====
```

✓CLI allows to retrieve values for all users ...

Cost output: eos inspector ls

```
Storage Costs – User View [ disk ]
```

```
-----  
Total Costs : 195.02 kEOS  
-----
```

```
01. atlas003           : 112.10 kEOS  
02. wguan              : 17.70 kEOS  
03. at11111           : 6737 EOS  
04. jt                 : 4213 EOS  
05. jb                 : 1736 EOS  
06. 40                 : 1683 EOS  
07. 77                 : 1519 EOS  
08. 16                 : 1213 EOS  
09. 16                 : 1177 EOS  
10. clscr              : 1137 EOS
```

- ✓CLI allows to retrieve values for all users ...
- ✓Currency can be changed to EUR, CHF, US\$, YEN

Cost output: eos inspector ls

```
=====
Storage Costs - User View [ disk ]
-----
```

```
Total Costs : 195.02 kEOS
-----
```

```
01. atlas003           : 112.10 kEOS
02. wguan              : 17.70 kEOS
03. atlas001          : 6737 EOS
04. jt                 : 4213 EOS
05. jb                 : 1736 EOS
06. 40                 : 1683 EOS
07. 77                 : 1519 EOS
08. 16                 : 1213 EOS
09. 16                 : 1177 EOS
10. clsr               : 1137 EOS
```

- ✓ CLI allows to retrieve values for all users ...
- ✓ Currency can be changed to EUR, CHF, US\$, YEN
- ✓ Disk & Tape PRICE/TB/YEAR can be configured or can be reported neutrally as TB/y

Summary & Outlook

- Devices interface gives you insight into your hardware and it's value

- Devices interface gives you insight into your hardware and it's value
 - registry allows to build additional functionality about hardware lifecycle

- Devices interface gives you insight into your hardware and it's value
 - registry allows to build additional functionality about hardware lifecycle
- Inspector interface provides a cost overview of data stored

- Devices interface gives you insight into your hardware and it's value
 - registry allows to build additional functionality about hardware lifecycle
- Inspector interface provides a cost overview of data stored
- **EOS** billing implementation

- Devices interface gives you insight into your hardware and it's value
 - registry allows to build additional functionality about hardware lifecycle
- Inspector interface provides a cost overview of data stored
- **EOS** billing implementation
- option could be to create monthly billing report per *user/group/app* about IO and meta-data usage, quota & occupancy

December 2023
Here is your monthly EOS invoice

# Files	1.8 M
# Directories	204k
Cost of your currently stored data	1.503 EOS
Cold data	89%
meta-data usage	1.3 M OPS
EGRESS	15.3 TB
INGRESS	3.5 TB

Advice: archive your cold files !

Sincerely CERN-IT Cost Working Group

- Devices interface gives you insight into your hardware and it's value
 - registry allows to build additional functionality about hardware lifecycle
- Inspector interface provides a cost overview of data stored
- **EOS** billing implementation
- option could be to create monthly billing report per *user/group/app* about IO and meta-data usage, quota & occupancy

Try it out!

December 2023
Here is your monthly EOS invoice

# Files	1.8 M
# Directories	204k
Cost of your currently stored data	1.503 EOS
Cold data	89%
meta-data usage	1.3 M OPS
EGRESS	15.3 TB
INGRESS	3.5 TB


Advice: archive your cold files !

Sincerely CERN-IT Cost Working Group




Useful Links

Web Page <https://eos.cern.ch>




GITLAB Repository <https://gitlab.cern.ch/dss/eos>

GITHUB Mirror <https://github.com/cern-eos/eos>



Community Forum <https://eos-community.web.cern.ch/>

email: eos-community@cern.ch



Documentation <http://eos-docs.web.cern.ch/eos-docs/>



Support email: eos-support@cern.ch

Thank you for your attention!
Questions?

