



# RNTuple Test in CMSSW: AOD Take 2

Dr Christopher Jones

CCE-SOP

09 October 2024

# Take 2

---

- Values reported last week were found to be made using an old version of ROOT
  - Pre dynamic page scaling
- Redid analysis using the correct version of ROOT

# AOD Format

---

- 2nd largest format that CMS produces
  - ~400kB/ event
- Stored to tap and kept on disk
- Characteristics
  - 552 top level branches
  - Fully split
  - LZMA level 4 compression

# RNTuple Storage of AOD

---

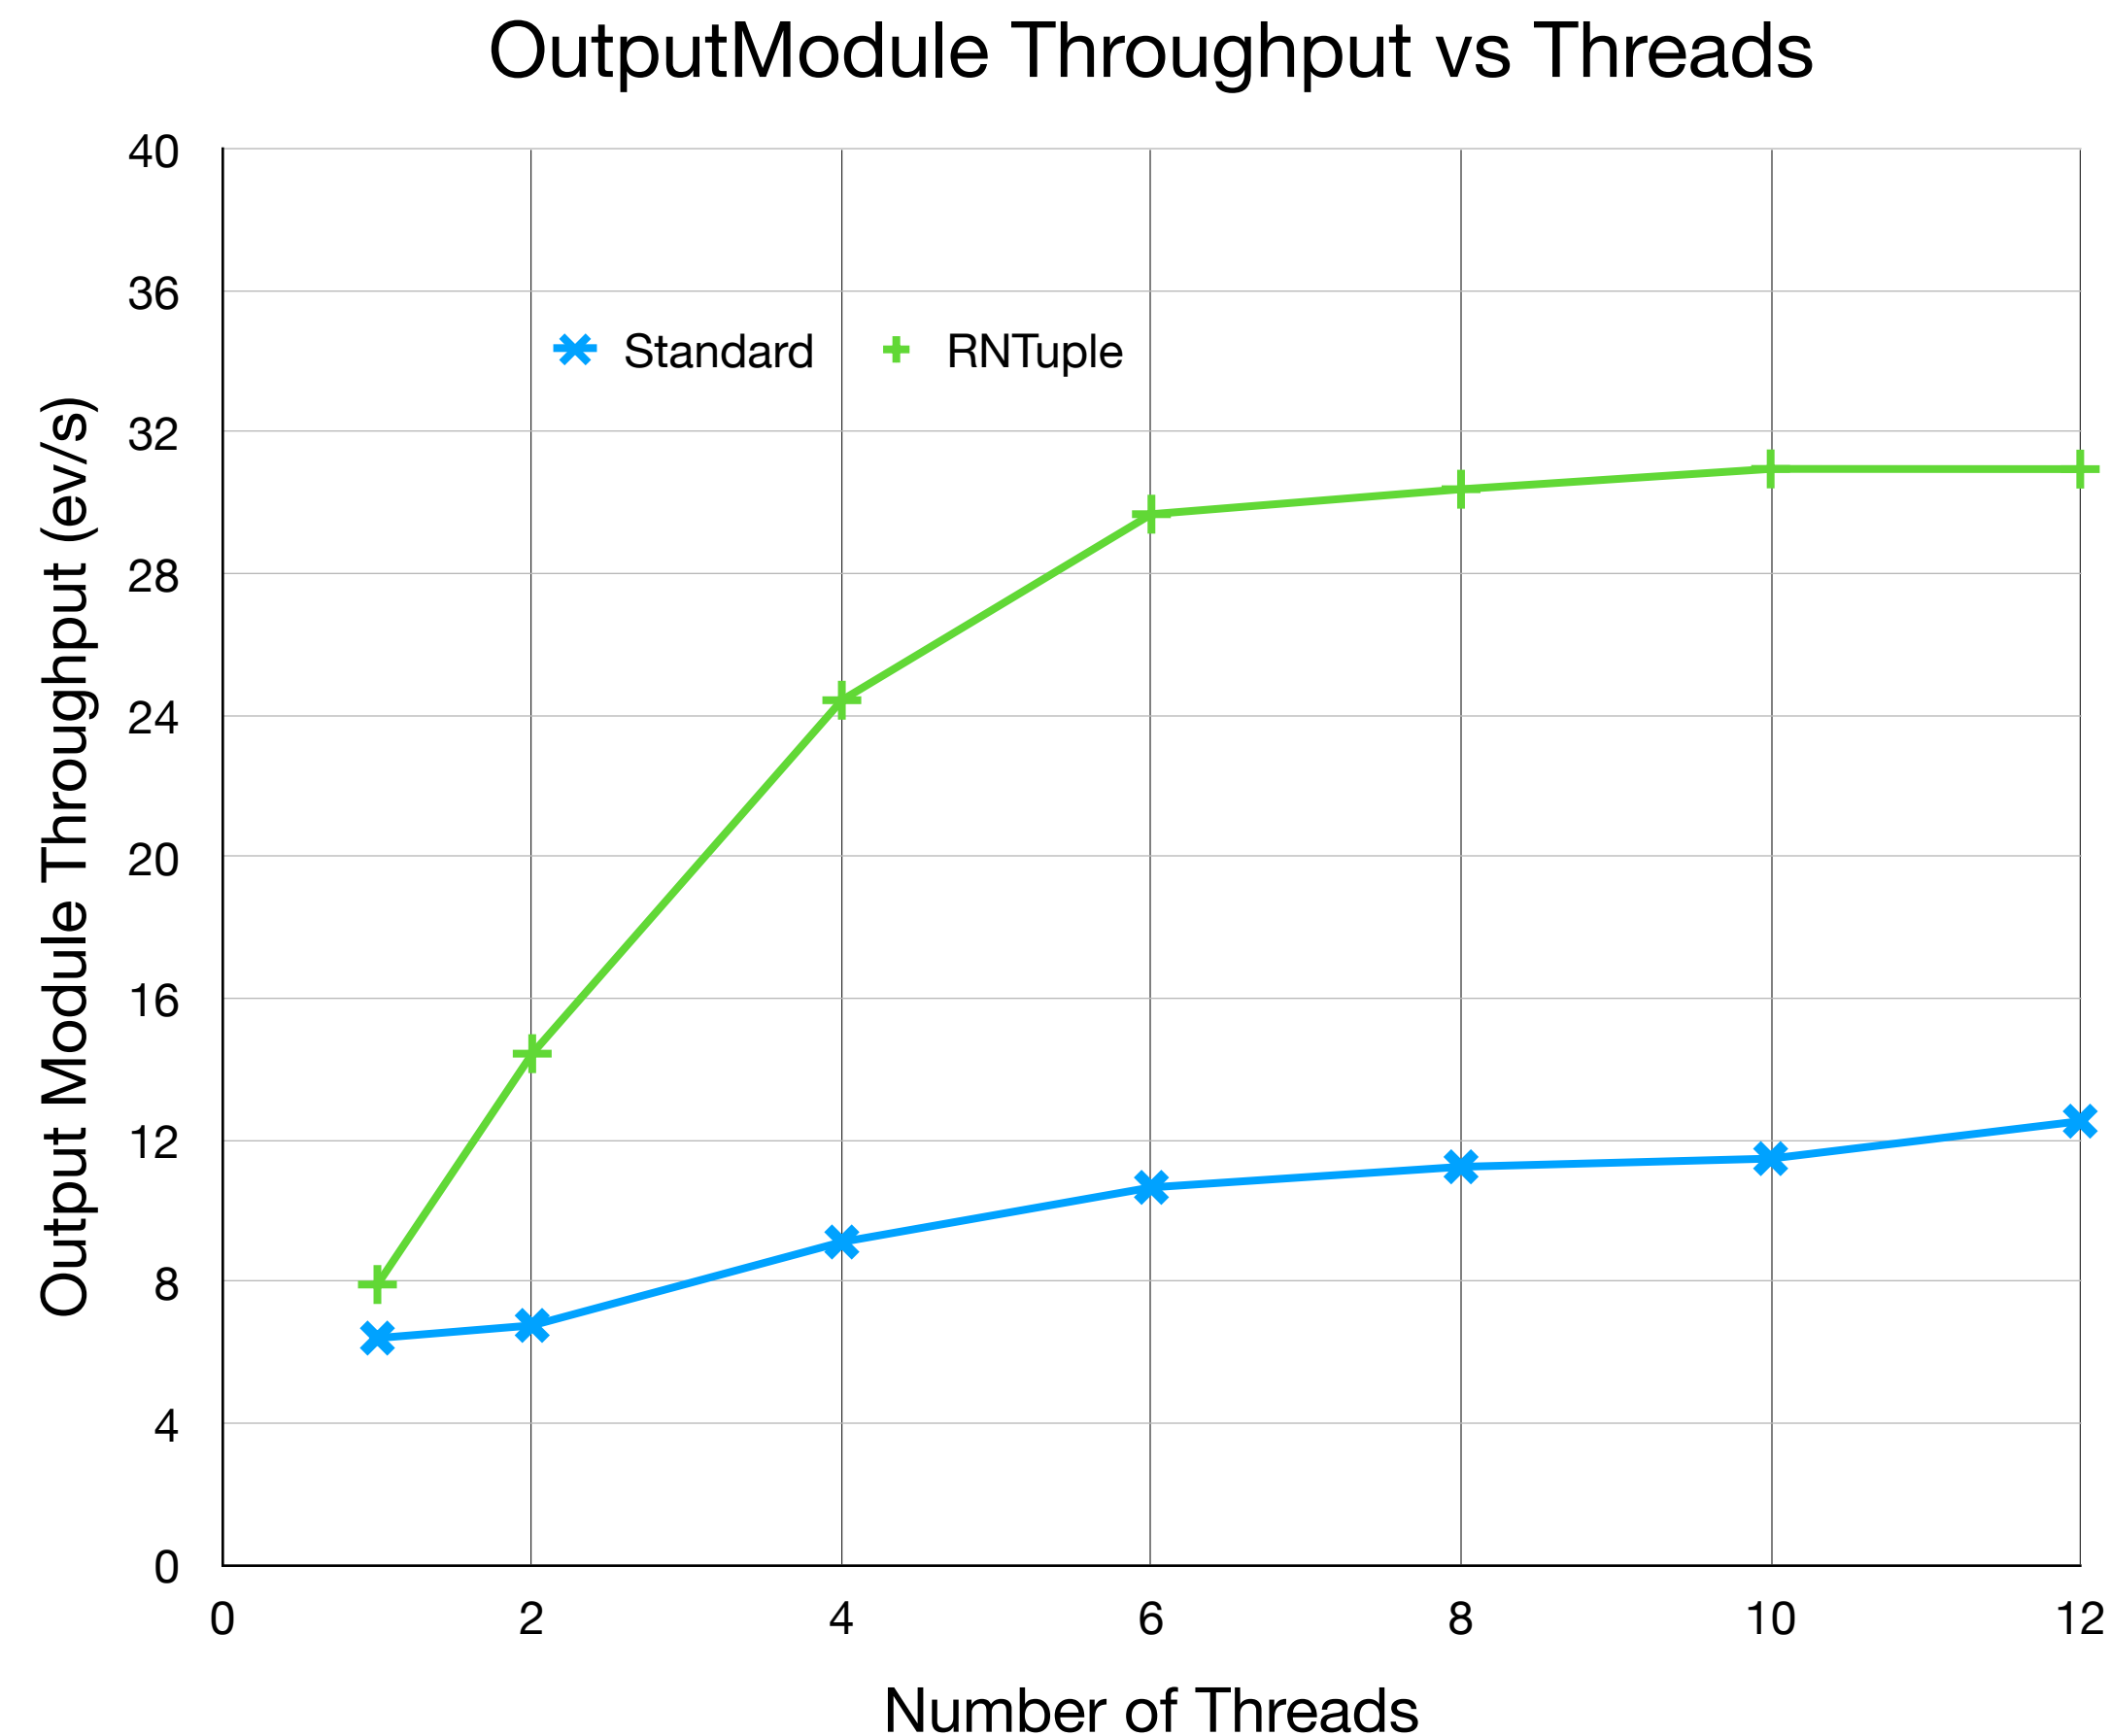
- Difference between fully split and partial split is negligible

File Sizes

File	Size	Relative %
TTree	4,000,822,391	100.00%
RNTuple	3,555,725,933	<b>88.87%</b>
RTuple No Split	3,938,704,279	98.45%
RNTuple Partial Split	3,522,148,557	88.04%

# Thread Scaling

- Measurement of only time spent doing output related work
- Good thread scaling compared to TTree



# Memory Usage

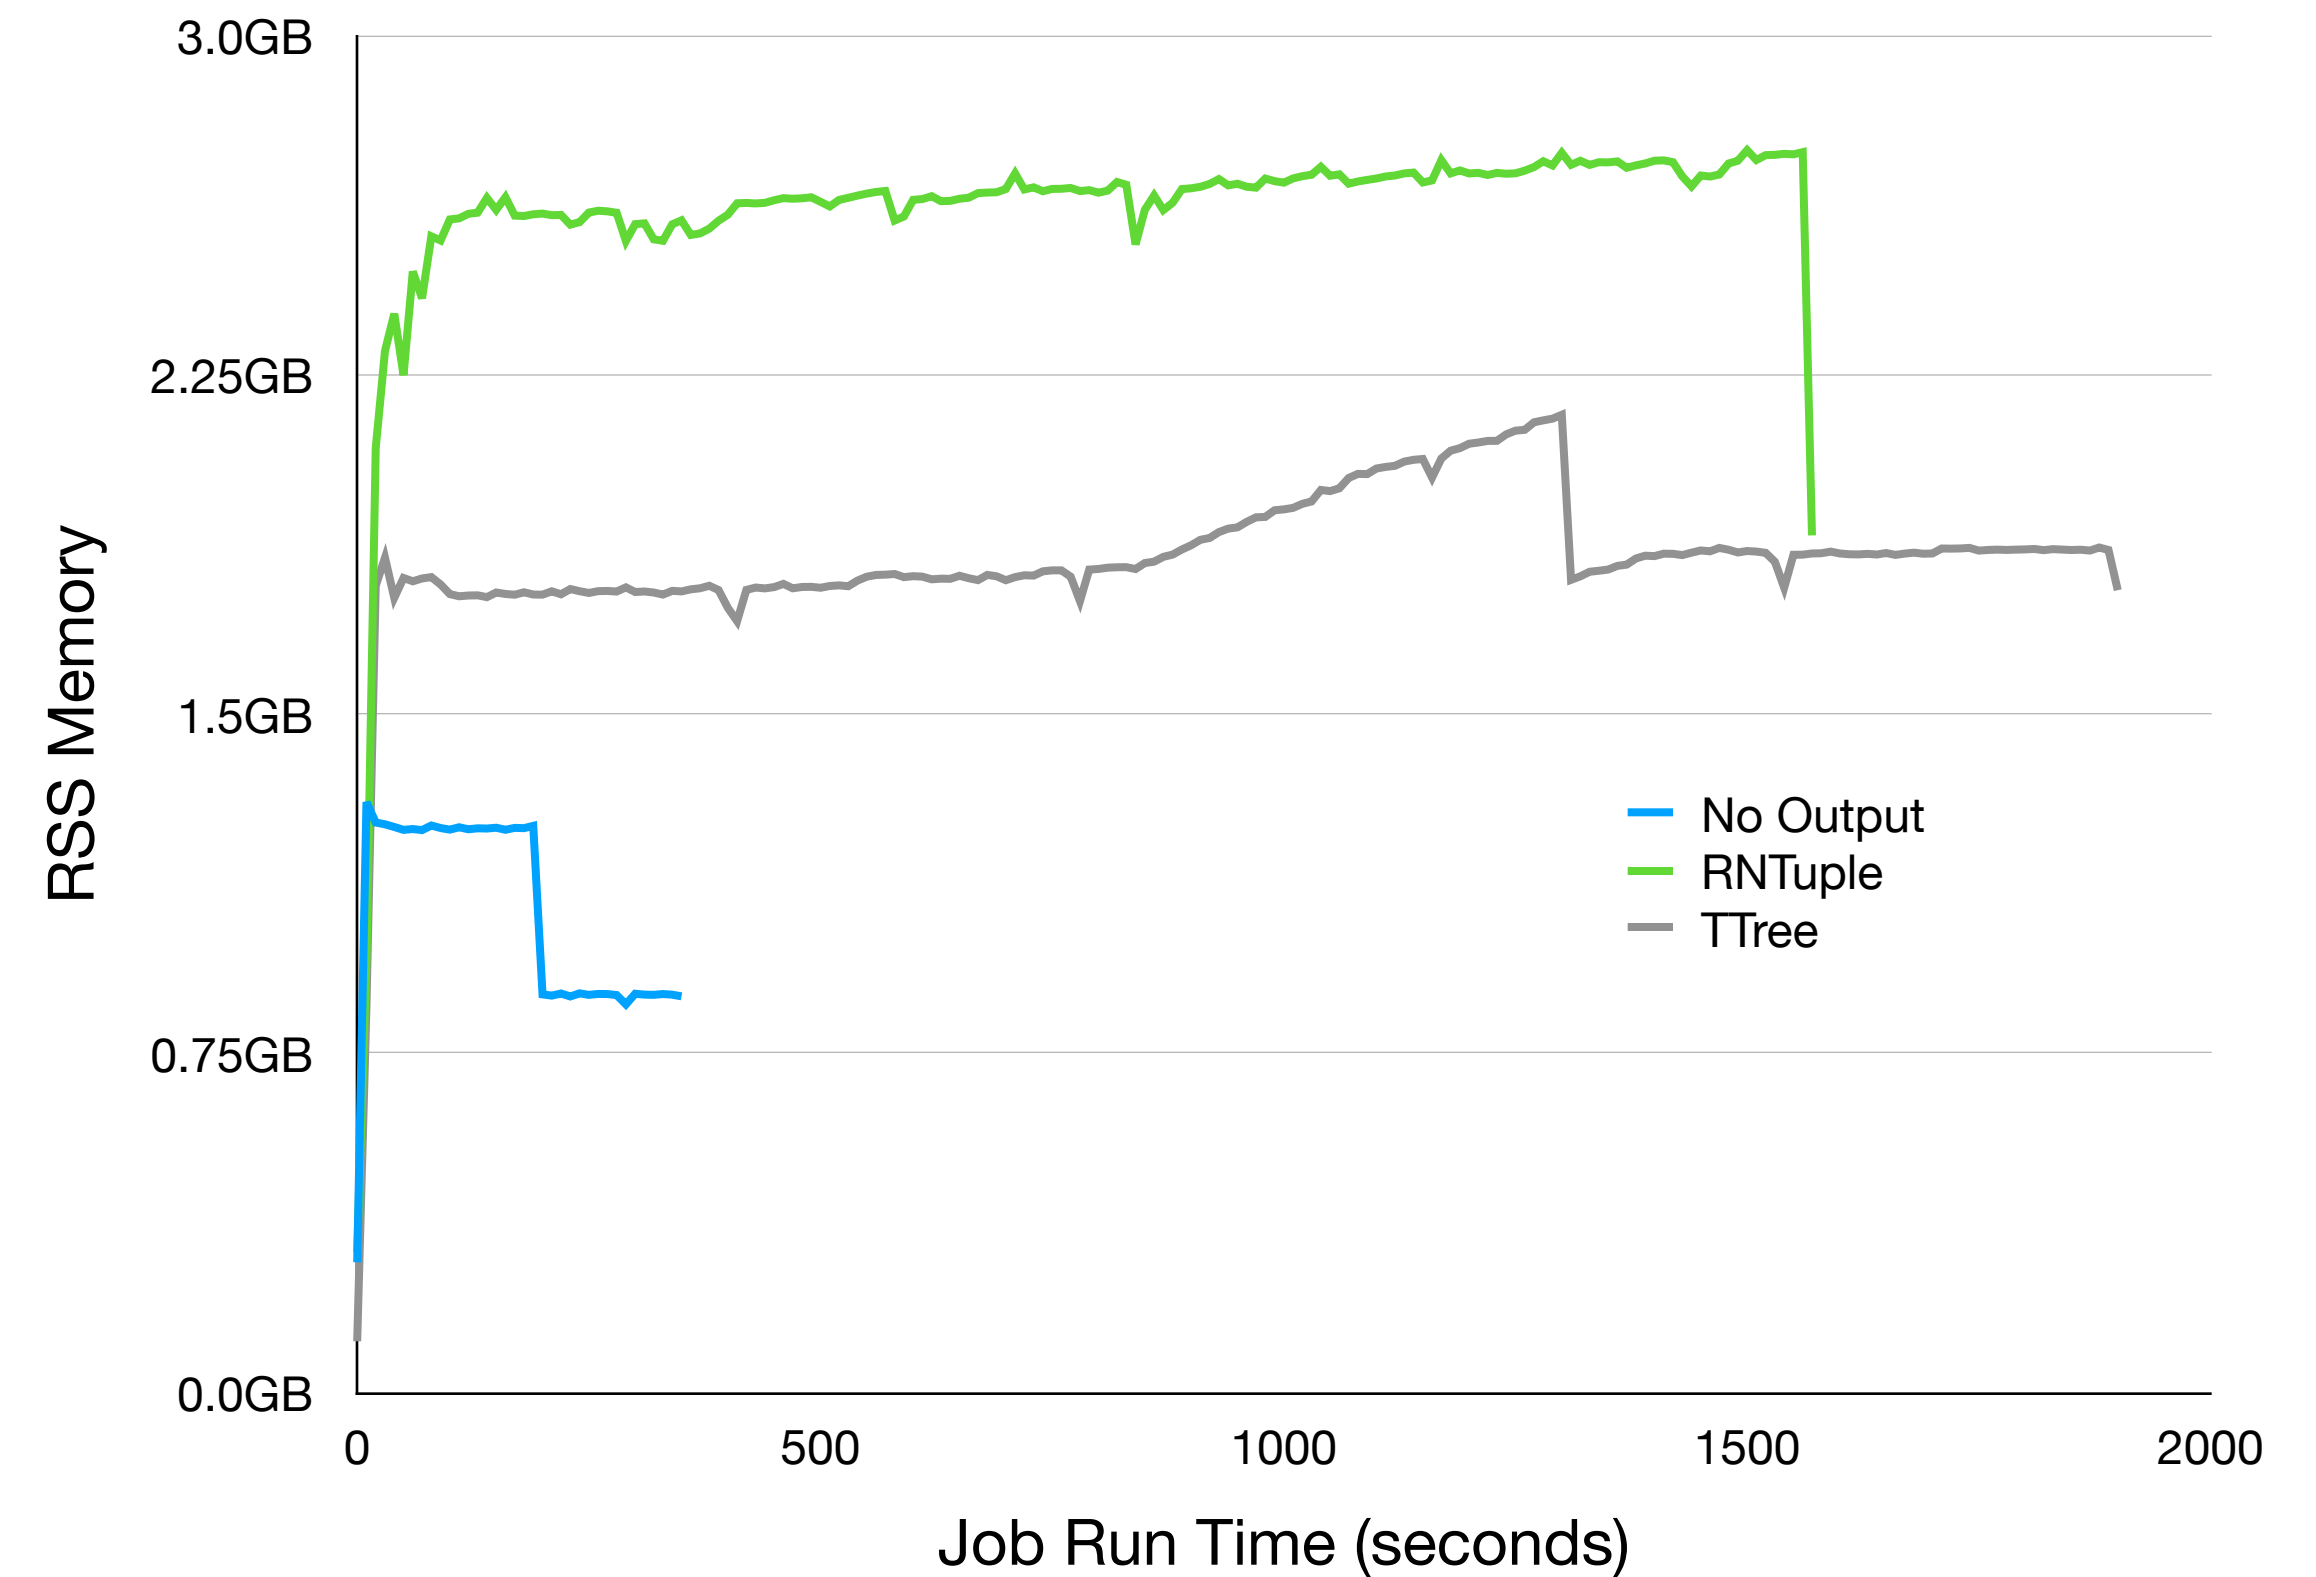
- Both maximum total requested allocations and RSS are poor

Max Total Allocation Requested

Job	Max Alloc (b)	Overhead	Ratio
No Output	842,128,400	0	
TTree	1,365,851,992	523,723,592	1.00
RNTuple	2,351,522,504	1,509,394,104	<b>2.88</b>

Max RSS

Job	Max RSS	Overhead	Ratio
No Output	1,336,922,112	0	
TTree	2,213,576,704	876,654,592	1.00
RNTuple	2,812,416,000	1,475,493,888	<b>1.68</b>



# RNTuple Buffered Writing

---

- RNTuple has an option to turn off buffered writing
- Buffered writing aids in optimizing file layout
  - puts related pages together on the file
  - decreases file size by  $\sim 2.5\%$

File Sizes

File	Size	Relative %
TTree	4,000,822,391	100.00%
RNTuple	3,555,725,933	<b>88.87%</b>
RTuple No Buffered Write	3,657,766,333	<b>91.43%</b>

# Memory Usage: No Buffered Write


- RNTuple has an option to turn off buffered writes
  - buffered writing helps with
- Memory usage vastly improved

Max Total Allocation Requested

Job	Max Alloc (b)	Overhead	Ratio
No Output	842,128,400	0	
TTree	1,365,851,992	523,723,592	1.00
RNTuple	2,351,522,504	1,509,394,104	2.88
RNTuple noBufWrite	1,543,954,624	701,826,224	<b>1.34</b>

Max RSS

Job	Max RSS	Overhead	Ratio
No Output	1,336,922,112	0	
TTree	2,213,576,704	876,654,592	1.00
RNTuple	2,812,416,000	1,475,493,888	1.68
RNTuple noBufWrite	2,228,609,024	891,686,912	<b>1.02</b>





# Thread Scaling : No Buffered Write

- Turning off buffered write turns off threading support

