

The Thrifty Admin

VDI by Day Compute by Night

Tony Foster – Principal Technical Marketing Engineer, Dell Technologies
VMware vExpert; VMware EUC Champion;
VMware Experts Program, BDSEW;
NVIDIA vGPU Community Advisor (NGCA)

@wonder_nerd www.wondernerd.net

V 1.0 Date:Aug 26, 2019

#VMworld #VMTN5019U. #vBrownBag #TechTalk



It's Coming...

That new project...


The executives want AI (they read about it in a trade magazine)...

You've been tasked with figuring out how to deliver it, but...

The budget is tight, most of it is going to a Data Scientists...

What's a thrifty admin to do?

Basics First




Created by Nhor
from Noun Project

Make a Plan



Created by Adrien Coquet
from Noun Project

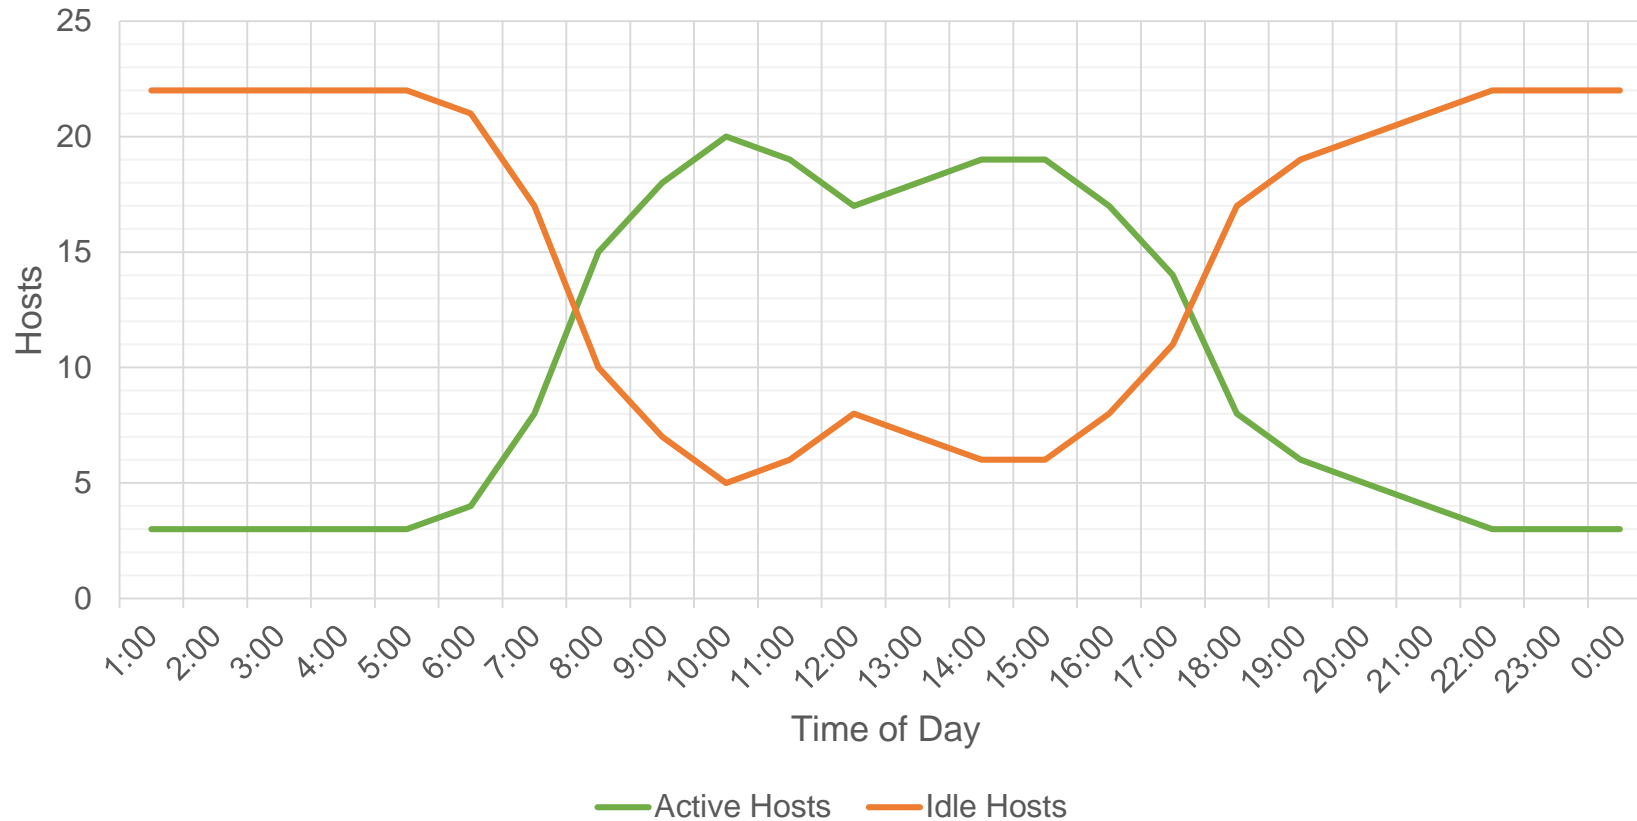
Build AI
Infrastructure



Enhance with
Spare Resources

What Spare Resources?

Daily VDI Usage
(M-F)



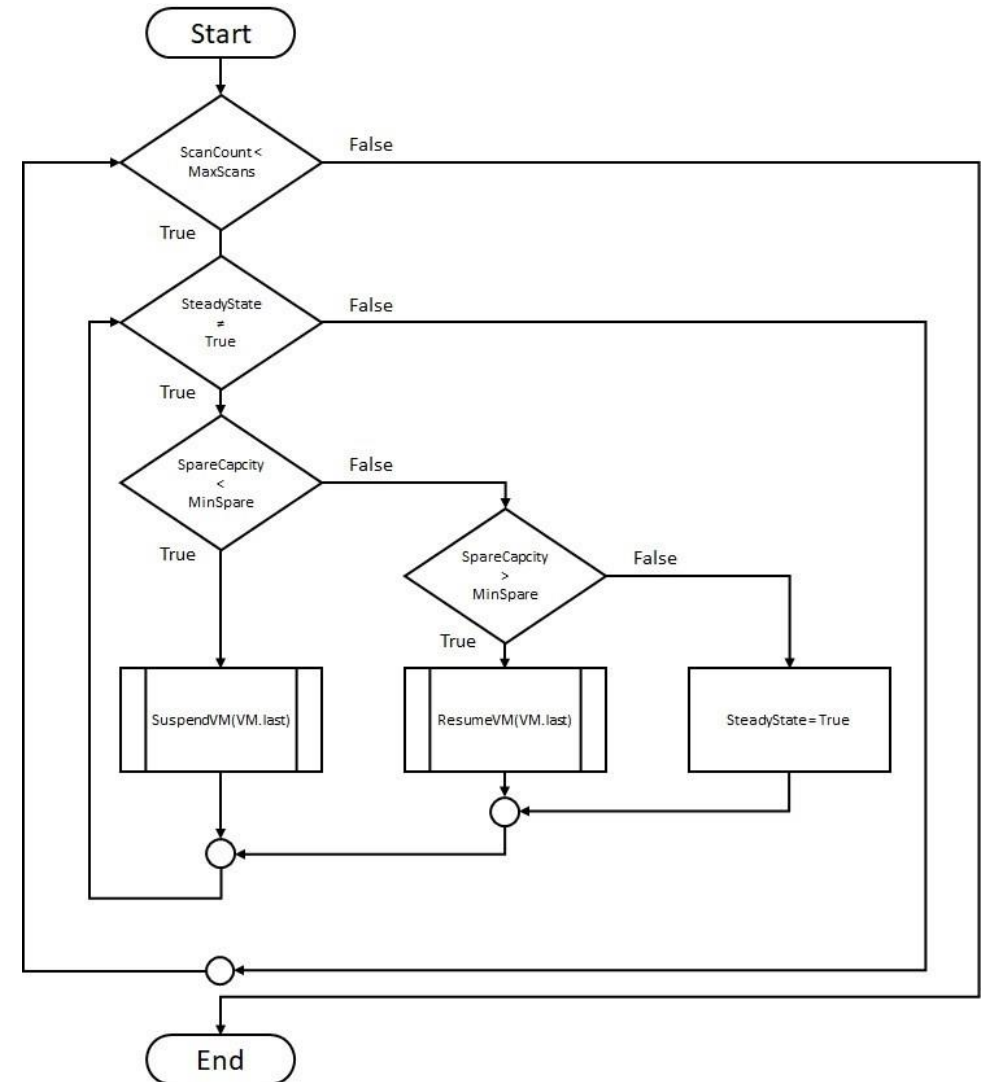
Time	Active Hosts	Idle Hosts	Idle GPUs
1:00	3	22	66
4:00	3	22	66
8:00	15	10	30
12:00	17	8	24
13:00	18	7	21
16:00	17	8	24
20:00	6	19	60
0:00	3	22	66

Harvesting Spare Resources

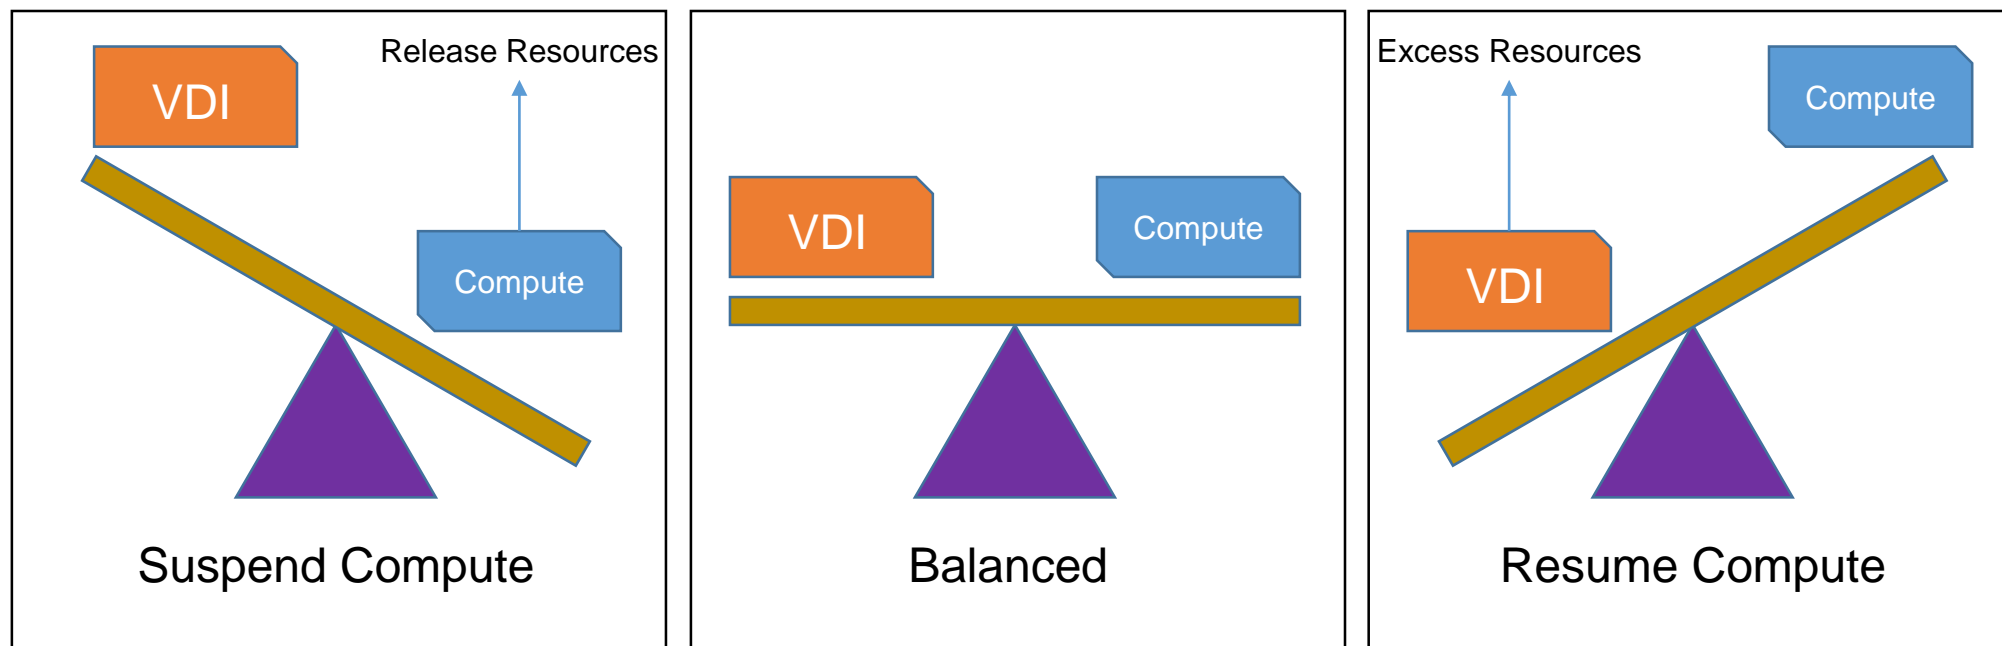
(VDI by Day Compute by Night)

Simple approach:

- If there are free resources use them for AI
Resume AI VM n
- If there are not enough resources for VDI Suspend
AI VM n
- Maintain a steady state
- Repeat

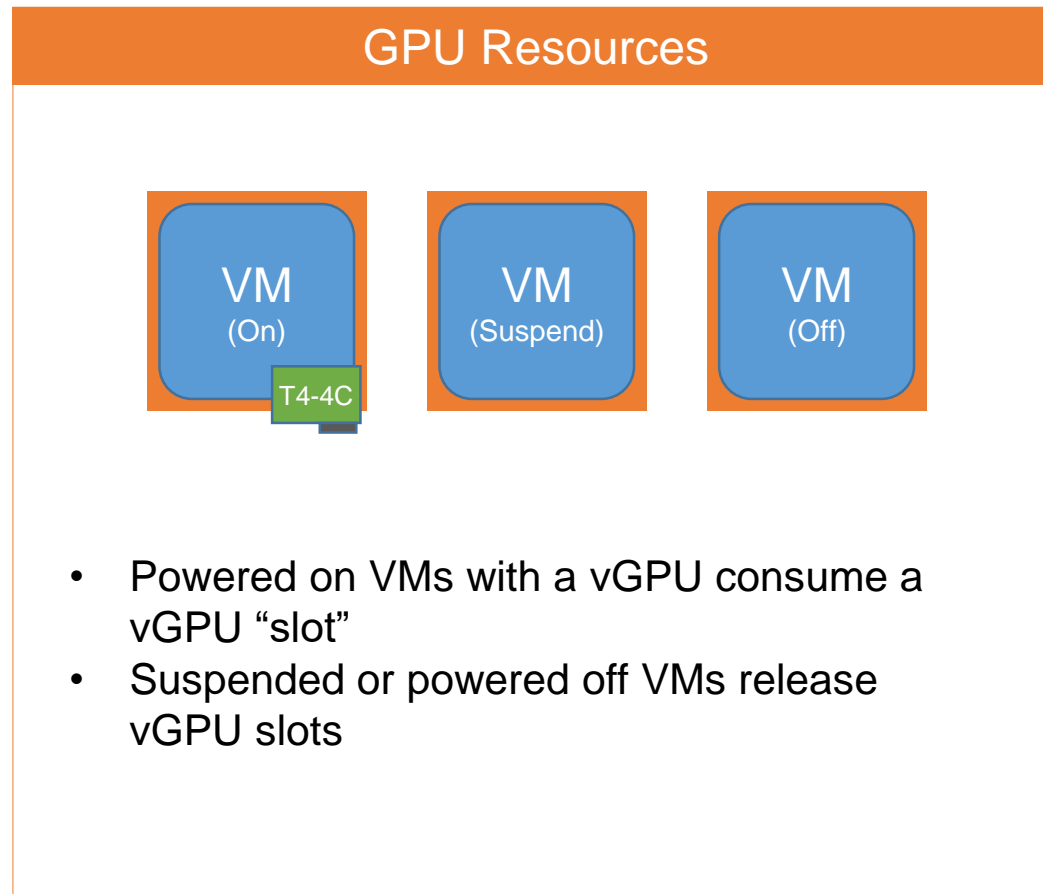
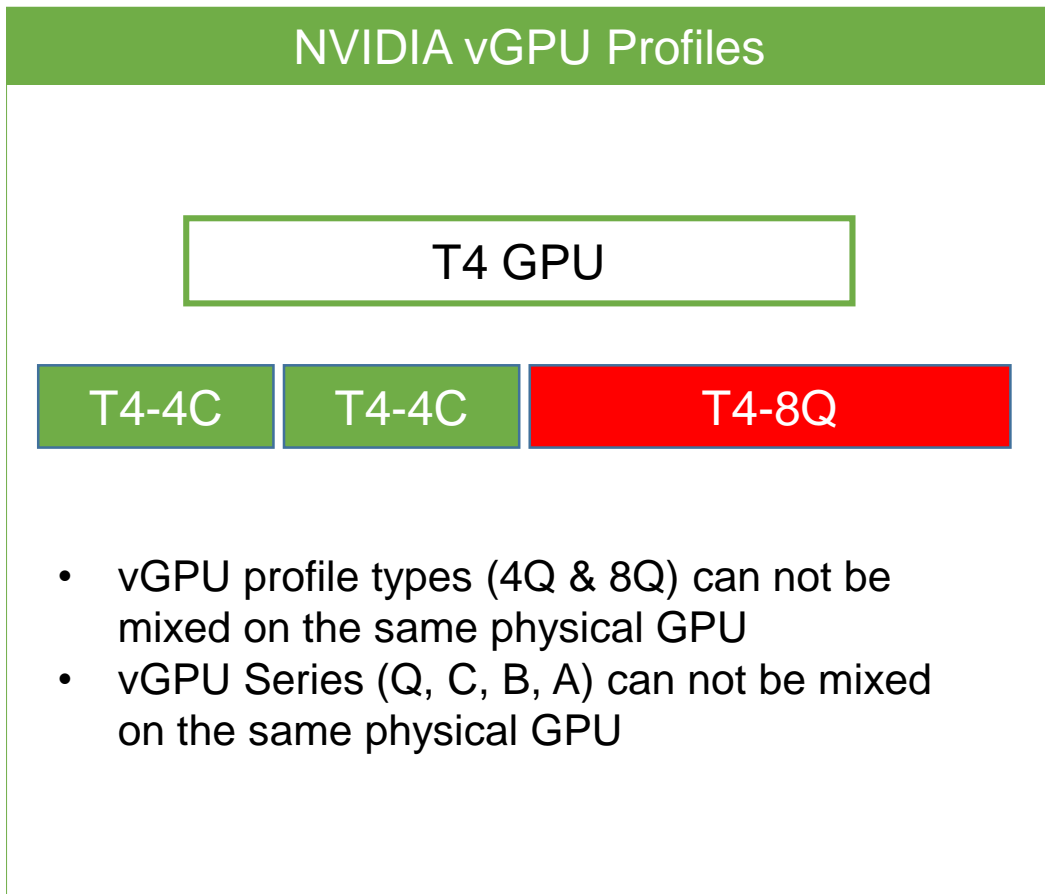


Something Like This...



That's Simple?

Sure, don't forget...



What Methods Can We Use?

SLURM



- <https://1drnrd.me/VDIbyDayCompute>
- Open-source workload manager designed for Linux clusters of all sizes
 - Allocates access to resources to users so they can perform work
 - Provides a framework for starting, executing, and monitoring work
 - Arbitrates contention for resources by managing a queue of pending work.

PowerCLI

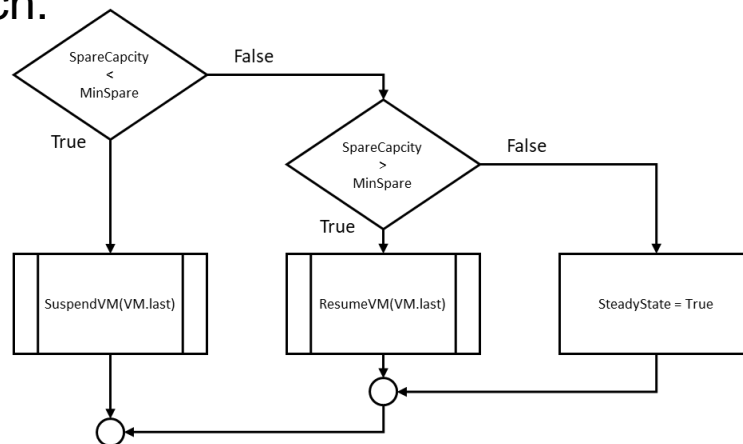


- VMware PowerCLI provides a PowerShell interface to the VMware product APIs
 - Based on Microsoft PowerShell
 - Uses the PowerShell basic syntax and concepts
- Most organizations and admins have experience with it

PowerCLI Approach

Control

- Suspends and Resume VMs
- vMotion VMs
- Approach:



Capacity

- Available vGPU slots
- vGPUs in use
- Approach:
 - Count GPUs in a Cluster (S)
 - Count vGPUs in use (U)
 - Find Profile Difference (D)
 - VDI Profile Size (V)
 - Compute Profile Size (C)

$$\text{Spare vGPU} = (S \times C) - ((S \times V) - U)$$

Sample Script and Details

Details

<https://www.wondererd.net/blog/scripting-vdi-by-day-and-compute-by-night>

<https://github.com/wondererd/VDIbyDayComputeOtherwise>

Calls

<https://github.com/wondererd/vGPUCapacity>

```
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
```

Don't Forget the Licensing

NVIDIA vGPU

- Different license types

Series	Optimal Workload
Q-series	Virtual workstations for creative and technical professionals who require the performance and features of Quadro technology
C-series	Compute-intensive server workloads, such as artificial intelligence (AI), deep learning, or high-performance computing (HPC)
B-series	Virtual desktops for business professionals and knowledge workers
A-series	App streaming or session-based solutions for virtual applications users

<https://docs.nvidia.com/grid/latest/grid-vgpu-user-guide/index.html#supported-gpus-grid-vgpu>

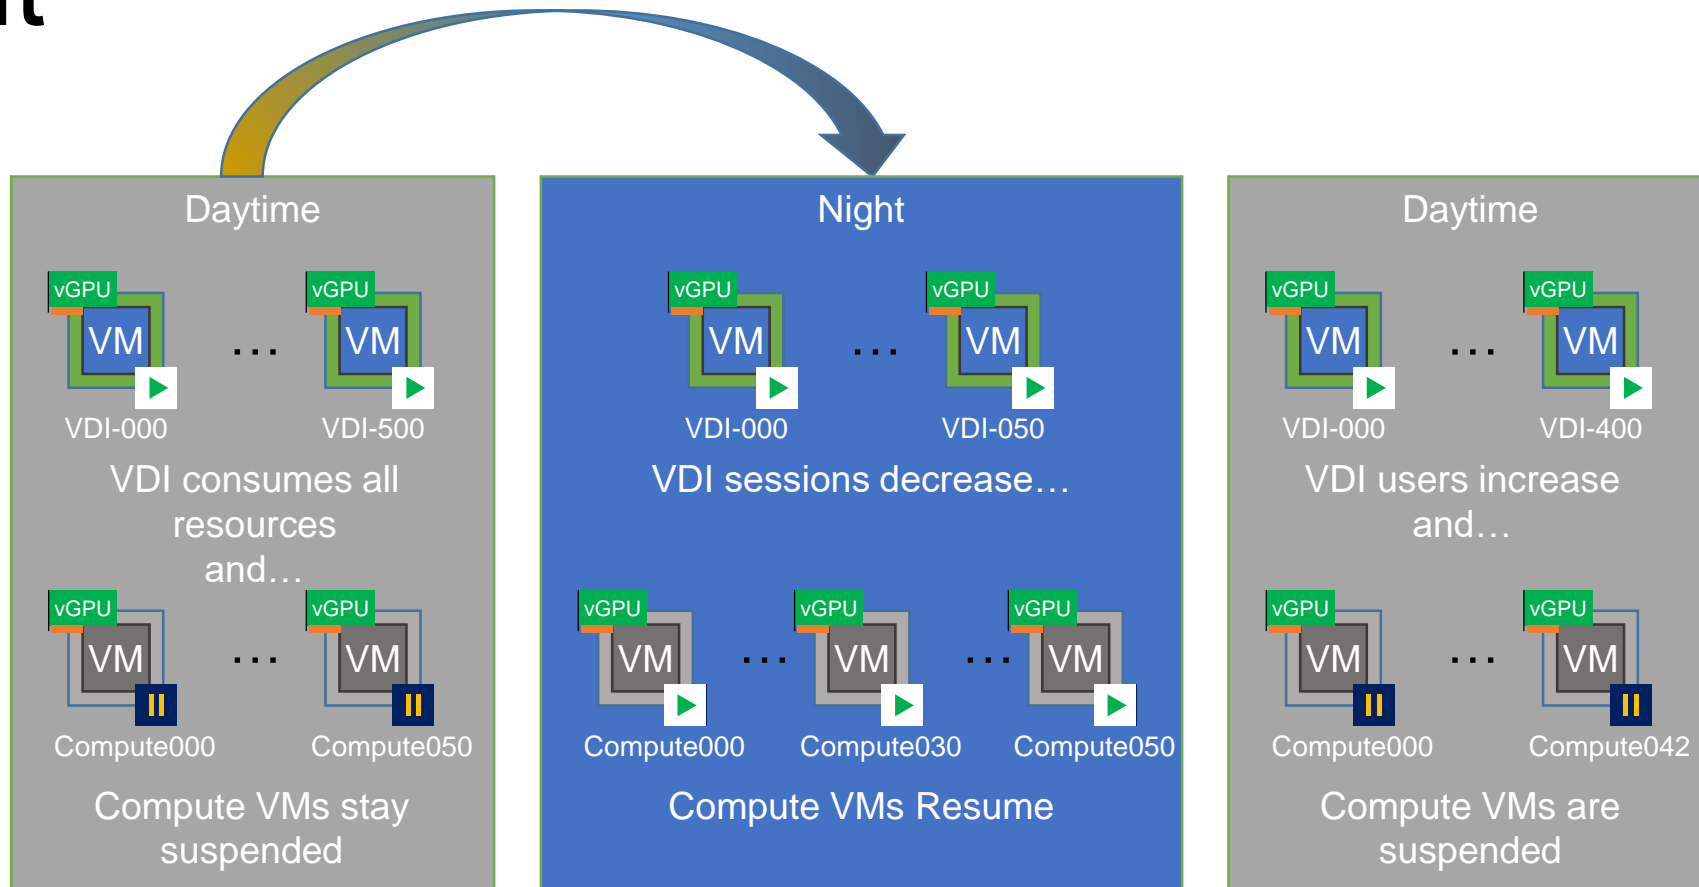
vSphere

VMware Horizon Licensing:

- “vSphere Desktop can be used to power virtual desktops, as well as the underlying infrastructure for the VDI environment”
- “For mixed workloads, we [VMware] recommend buying an edition of vSphere, which is licensed per CPU and buying a Horizon Add-on to run on top of that”
- VMware vSphere Scale Out Edition is one option

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/whitepaper/vm-w-workspace-one-horizon-uselet.pdf>


Giving us VDI by Day and Compute By Night



Catch These VMworld Sessions

- **HBI1546BU**
How GPU-Assisted ML for Medical Research Proved to Be a Force for Good
Johan van Amersfoort and Niels Hagoort
Tuesday 11:00 AM
- **MLA3388BU**
Scaling the Virtualization Hurdle for Machine Learning, Big Data and HPC
Aviad Shaul Yehezkel and Adit Ranadive
Wednesday 8:00 AM
- **MLA3014WU**
Supporting Machine Learning Workloads and GPUs on vSphere
Justin Murray and Josh Simons
Thursday 10:30AM
- **MLA3390BU**
Running GPU-Accelerated Data Science Workflows Virtually using NVIDIA vGPU
Raj Rao and Manvender Rawat
Thursday 12:00 PM

Questions



Thank you for attending this
vBrownBag Tech Talk and supporting the VMTN
Slides are available at <https://1drnrd.me/VMTN5019U>

Tony Foster
@wonder_nerd
Tony.Foster@wondernerd.net
<https://wondernerd.net>