

VMware Horizon and PowerVault ME4 Storage White Paper

A technologies overview of ME4 being used in small VDI deployments of up to 50 users

Acknowledgements

This whitepaper was made possible thanks to the contributions from:

- Tim Carlson
- Joe Catalanotti
- Daniel Jobke
- Martin Pritchard
- Selim Selveroglu

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © February 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [7/28/2020] [Whitepaper] [VMware Horizon and PowerVault ME4 Storage White Paper]

Contents

| | | |
|-----|---|-------------------------------------|
| 1 | Introduction..... | 3 |
| 2 | Physical Setup..... | Error! Bookmark not defined. |
| 2.1 | Deployed Hardware | 4 |
| 2.2 | Cabling Diagram..... | 5 |
| 3 | Software and Licence Requirements | 6 |
| 3.1 | Required Software..... | 6 |
| 3.2 | Required Software Licences | 6 |
| 4 | Environment VMs | 7 |
| 4.1 | VMware | 7 |
| 4.2 | Microsoft..... | 7 |
| 5 | VMware Horizon Pool Configuration | 9 |
| 6 | VMware Host configuration | 10 |
| 7 | Performance | 11 |
| 7.1 | Testing Methods..... | 11 |
| 7.2 | Testing Results..... | 13 |
| 8 | Conclusion..... | 16 |

Introduction

The scope of this solution is to ascertain what kind of workloads and performance can be achieved. For this purpose, the ME4 with iSCSI controllers were chosen, as it leaves users the capability to grow their environment beyond the 4-host direct-attached limit of the SAS controllers by introducing switches, as well as, the option to start storage-based replication down the line.

If combined with VRTX server chassis this would enable us to sell a high-performance and cost-effective solution in a 7U formfactor, which can easily be upscaled on the storage and server end.

In our configuration, the virtual machines are not using persistent disks. If a user logs out of Windows, the Virtual Machine gets deleted and re-created from the template VM. The expectation is that users store their personal or work data on personal/departmental shares or SharePoints. Furthermore, we are looking at a single virtual pool configuration, which enables the user to easily use the second virtual pool to store their applications like Exchange, File, Database servers or infrastructure services, such as Active Directory Domain Controllers, DNS, DHCP servers.

For VMware Horizon to properly work a machine (virtual or physical) with DHCP and DNS role is required. Failing to have a DHCP address scope adequate to sustain the desired number of users/VDI instances, will result in VMware Horizon Connection Server being unable to present available machines to users, therefore machine will show in error state in the VMware Horizon View Connection Manager Administrator Console.

2. Physical Setup

2.1 Deployed Hardware

Server Hardware:

R730 - ME4SRVR01.me4.test:

- Dual Xeon E5-2687W (10 Cores @ 3.1Ghz)
- 256GB DDR4-2133
- LOM: Dual Intel X540 & Dual I350
- Local storage: dual 15K 300GB SAS drives in R1
- Dual SD card boot device for ESXi 6.7U3

R730 - ME4SRVR02.me4.test:

- Dual Xeon E5-2697v3 (14 Cores @ 2.6Ghz)
- 256GB DDR4-2133
- LOM: Dual Intel X540 & Dual I350
- Local Storage: 6x 15K 300GB SAS drives in R5
- Dual SD card boot device for ESXi 6.7U3

Switch:

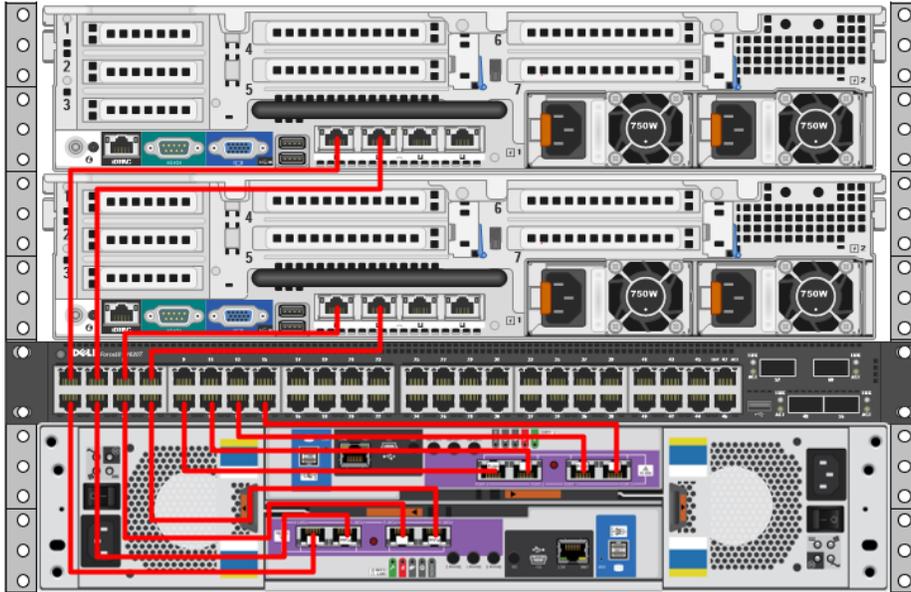
Force10 S4820T

Storage:

ME4024iSCSI

- Dual Controller
- 4x 10GBps Ports per controller
- 2x 1.92 TB SSDs in R1
- 10x 1.2TB 10K SAS in R6
- Virtual Mode with single pool
- Single 10TB Volume spanning the pool

2.2. Cabling Diagram



For the purpose of this testing, we only attached the ME4024 iSCSI to a single S4810 switch. In a production environment we strongly encourage user to either be attached to 2 switches or direct attaching their hosts to the ME4.

3. Software and Licence Requirements

3.1 Required Software

VMware:

- VMware vSphere Server 6.7U3
- VMware vSphere Hypervisor (ESXi) 6.7U3
 - vDRS fully automated enabled
- VMware Horizon 7.12.0
 - View Agent 64bit
 - View Agent Direct-Connection 64bit
 - Persona Management 64bit
 - View Connection Server 64bit

Microsoft:

- Windows Server 2016 DC
 - Active Directory Domain Controller
 - DNS
 - DHCP
- Windows 10 Enterprise
 - 7ZIP
 - Google Chrome
 - Adobe Acrobat
 - Windows Office

PowerVault ME4:

- G280R004-01

3.2 Required Software Licences

VMware:

- VMware vSphere Server Standard
- VMware vSphere Enterprise Plus
- VMware Horizon Trial Licence (Enterprise)

Microsoft:

- Windows Server 2016 DC
- Windows 10 Enterprise

4. Environment VMs

4.1 VMware

| | |
|------------------------------|---------------------------|
| vSphere Server (VACS) | |
| vCPUs | 2 |
| Memory | 10GB |
| Storage | 38GB stored on ME4 volume |
| Version | 6.7U3 |

4.2 Microsoft

| | |
|------------------------------|--|
| Horizon Management VM | |
| vCPUs | 6 |
| Memory | 16GB |
| Storage | 60GB stored on ME4 volume |
| OS | Microsoft Server 2016 DC |
| Horizon Software | View Connection Sever 7.12.0 64bit Persona Management 7.12.0 64bit HTML Access 5.4.0 64bit |

| | |
|-------------------------------------|---|
| Windows Management VM | |
| vCPUs | 6 |
| Memory | 16GB |
| Storage | 50GB stored on ME4 volume |
| OS | Microsoft Server 2016 DC |
| Microsoft Roles and Features | Active Directory Domain Controller DNS DHCP |

| | |
|----------------------------|--|
| Horizon Template VM | |
| vCPU | 2 |
| Memory | 4 |
| Storage | 40GB stored on ME4 volume |
| OS | Windows 10 Enterprise 64bit |
| Horizon Software | View Agent 7.12.0 View Agent Direct-Connection 7.12.0 |
| Additional Software | 7ZIP V19.00 Adobe Acrobat V20.006.20042 Google Chrome V81.0.4044.113 |

Connection Server: The Connection Server is the central hub for managing connections. Users connect using either the Horizon View client or a web browser. The Connection Server then connects the user to the proper virtual desktop session based on the assignments. Users can either be assigned a static machine (in this case, for persistence) or given a random machine in a non-persistent environment.

Horizon View agent: Installed on user devices, the Horizon View agent provides users with quick, secure, multi-protocol access to documents, applications, and desktops from any of the user's devices including smartphones, tablets, and PCs.

Management console: The Horizon View management console is a web service application that can be run from any Windows server in the environment.

5. VMware Horizon Pool Configuration

| Type | Automated Desktop Pool |
|--------------------------------|---|
| vCenter Server | Instant Clone |
| User Assignment | Dedicated with automatic assignment |
| Storage Optimization | VSAN not used / unavailable |
| Desktop Pool ID | Horizon-Win-Pool |
| Provisioning Settings | <u>Basic:</u> Enable Provisioning Stop Provisioning on Error <u>Virtual Machine Name:</u> Horizon-Win <u>Max Number of Machines:</u> 30 |
| vCenter Settings: | <u>Parent VM:</u> /DC1/vm/Horizon/Win10-Horizon <u>Snapshot (needs to be taken offline):</u> /VM Snapshot 4%25252f17%25252f2020, 9:10:17 AM <u>VM Folder Location:</u> /DC1/vm/Horizon/Horizon-Win-Pool <u>Cluster:</u> /DC1/host/Cluster1 <u>Resource Pool:</u> /DC1/host/Cluster1/Resources <u>Datastores:</u> Horizon-DS1 (hosted by PowerVault ME4) |
| Desktop Pool Settings | Mostly left default, except: Enabled Reclaim VM disk space Initiate reclamation @ 5GB |
| Remote Display Settings | Mostly left default, except: HTML Access enabled |
| Guest Customization | Left at default |

Instant Clone: An instant-clone desktop pool is an automated desktop pool. vCenter Server creates the desktop VMs based on the settings that you specify when you create the pool.

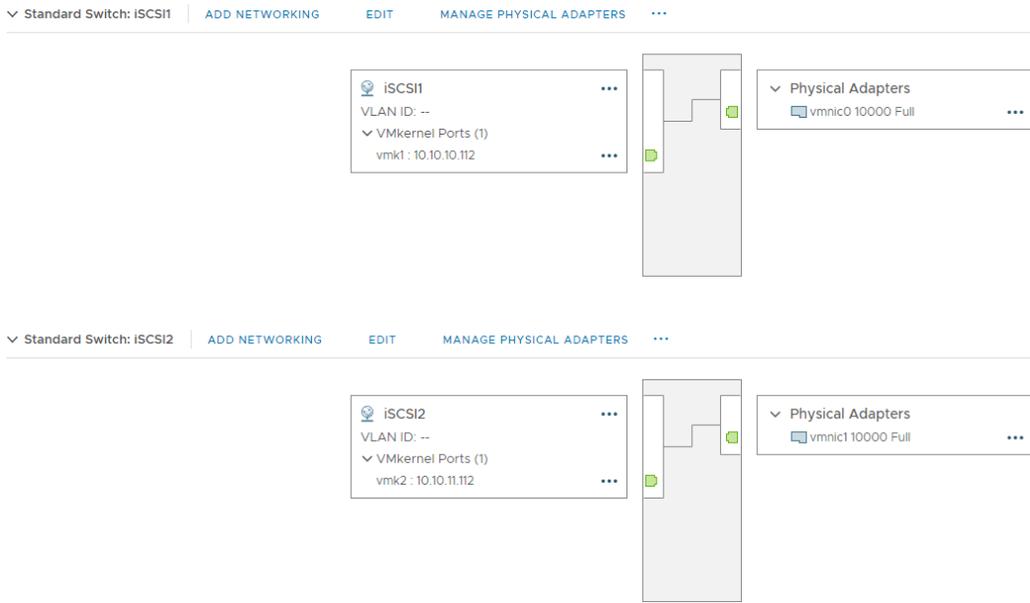
Similarly, to View Composer linked clones, an instant-clone desktop pool is based on a parent VM in vCenter Server, known as the master image. For instant-clone desktops, a parent VM is an internal VM that Horizon 7 creates and maintains, which is based on the master image. You cannot modify this internal parent VM. However, you can make changes to the master image.

Instant clones share a virtual disk of the master image and therefore consume less storage than full VMs. In addition, instant clones share the memory of the master image. Instant clones are created using the vmFork technology. See [>here<](#).

6. VMware Host configuration

Both physical hosts were configured in a fully automated vDRS cluster. This was done to help us load balancing the hosts. Each host was configured as per “[Dell EMC PowerVault ME4 Series and VMware vSphere](#)” guide, with 2 standard vSwitches, one per subnet, the MTU set to 8860 (both on the vSwitch and the VMKernel port) and delayed ACK disabled in the advanced initiator settings.

Standard vSwitches:



MPIO for the volume is set to Path selection policy: VMW_PSP_RR with 3 IOPS:

```
naa.600c0ff0003c94385429985e01000000
Device Display Name: DellEMC iSCSI Disk (naa.600c0ff0003c94385429985e01000000)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config: {implicit_support=on; explicit_support=off; explicit_allow=on; alua_followover=on; action_OnRetryErrors=on; {TPG_id=0,TPG_state=AO}{TPG_id=1,TPG_state=ANO}}
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config: {policy=rr, Iops=3, bytes=10485760, useANO=0; lastPathIndex=2: NumIOsPending=0, numBytesPending=0}
Path Selection Policy Device Custom Config:
Working Paths: vmhba64:C0:T0:L0, vmhba64:C6:T0:L0, vmhba64:C4:T0:L0, vmhba64:C2:T0:L0
Is USB: false
```

7. Performance

7.1 Testing Methods

Generating real world workloads tends to be a tricky endeavour. To measure what the ME4 and the hosts are theoretically capable of in this setup, the following 2 stages were performance tested:

1st stage:

50 dummy user accounts were created with same password for easier login. The 50 users were added to the Domain Admin security group, to allow diskspd test in stage 2. The Horizon View Client was installed on an

```
# Change working directory to view client installation folder
cd 'C:\Program Files (x86)\VMware\VMware Horizon View Client'
# Simple for-loop that declares i = 1, checks i is < 51 and then increases it by 1
for ($i=1; $i -lt 51; $i++){
    # start the actual session, substituting the $i for its current value of $i
    .\vmware-view.exe -serverURL horizon.me4.test -username user$i -password
    Password01 -domainName ME4 -desktopName Horizon-Win-Pool -standAlone -
    logInAsCurrentUser False -nonInteractive
}
```

external physical machine with access to the virtual machine network. The following PowerShell for-loop was used to launch the sessions: This will start 50 sessions in rapid succession, nearly instantaneous creating a login storm on the ME4.

2nd Stage:

Once the 50 sessions are started and the dummy users are logged in, after another 15sec Windows Task Scheduler will start a second PowerShell script, which will:

- Start Internet Explorer
- Start Outlook
- Start Word
- Run 10min diskspd test with 10MB file size & 4K block size - random
- Run 10min diskspd test with 10MB file size & 512 block size – random
- Run 10min diskspd test with 2GB file size & 512 block size - sequential

```

# wait 20s to allow user login to finish, then start IE
start-sleep -s 20
start-process -FilePath "C:\Program Files\Internet Explorer\iexplore.exe"
# wait 90s to simulate user logging into time management, then start outlook
start-sleep -s 90
Start-Process -FilePath "C:\Program Files (x86)\Microsoft Office\Office15\OUTLOOK.EXE"
# wait 10min to simulate user working in outlook, then start word
start-sleep -s 600
Start-Process -FilePath "C:\Program Files (x86)\Microsoft Office\Office15\WINWORD.EXE"

# 10 min for IOs to settle, then start diskspd tests
start-sleep -s 600
cd "C:\"
# first test small file with small blocksize to get max IOPs
.\diskspd.exe -c10M -d600 -L -r -Sh -o8 -t4 -a -w28 -b4K C:\iotest1.txt >
C:\iotest_10M_4K.txt
start-sleep -s 60
# second test small file with larger blocksize to get best throughput
.\diskspd.exe -c10M -d600 -L -r -Sh -o8 -t4 -a -w28 -b512k C:\iotest1.txt >
C:\iotest_10M_512.txt
start-sleep -s 60
# sequential test to observe tiering
.\diskspd.exe -c2G -d600 -L -si -Sh -o8 -t4 -a -w28 -b512k C:\iotest2.txt >
C:\iotest_2G_512.txt

```

Above the PowerShell script used to generate a workload on the virtual machines.

Arguably a “normal” user will never be able to saturate the virtual machine as much as diskspd is capable of. The diskspd test should be seen as worst-case scenario. Initially the sequential workload was generated by copying a 2GB file from a fileshare to the VM. Interestingly, such workload ended-up being perceived as random on the ME4, resulting in the spinning disks not being utilized at all.

7.2. Testing Results

Following the test results from the 2 staged benchmarking. IOPS are taken from the host-ports on the ME4 unless otherwise indicated.

Stage 1 – Login-storm:

On the Host ports we observed:

25825 IOPS
473 MB/s

While on the disks only:

| Tier I/O Information | | | | | | |
|----------------------|-------------|------------|----------|----------|-------|-------|
| Tier | A | | B | | Total | |
| | IOPS | MB/s | IOPS | MB/s | %IOPS | %MB/s |
| Performance | 7430 | 137 | | | 100% | 100% |
| Standard | 0 | 0 | | | 0% | 0% |
| Total | 7430 | 137 | 0 | 0 | | |

Stage 2 – IO-Storm:

Most IOPS observed while launching applications:

Host Ports:

22122 IOPS
952 MB/s

Disks:

| Tier I/O Information | | | | | | |
|----------------------|--------------|------------|----------|----------|-------|-------|
| Tier | A | | B | | Total | |
| | IOPS | MB/s | IOPS | MB/s | %IOPS | %MB/s |
| Performance | 11099 | 630 | | | 100% | 100% |
| Standard | 0 | 0 | | | 0% | 0% |
| Total | 11099 | 630 | 0 | 0 | | |

Diskspd run 1 – 10MB @ 4K:

Host Ports:

138352 IOPS
571 MB/s

Disks:

| Tier I/O Information | | | | | | |
|----------------------|--------------|------------|----------|----------|-------|-------|
| Tier | A | | B | | Total | |
| | IOPS | MB/s | IOPS | MB/s | %IOPS | %MB/s |
| Performance | 18746 | 588 | | | 100% | 100% |
| Standard | 0 | 0 | | | 0% | 0% |
| Total | 18746 | 588 | 0 | 0 | | |

Total Results from within a VM:

Total IO

thread | bytes | I/Os | MiB/s | I/O per s | AvgLat | LatStdDev | file

| | | | | | | | |
|--------|------------|---------|-------|---------|--------|-------|---------------------------|
| 0 | 1903943680 | 464830 | 3.03 | 774.70 | 10.324 | 4.613 | C:\iotest1.txt (10240KiB) |
| 1 | 1444122624 | 352569 | 2.30 | 587.60 | 13.613 | 7.127 | C:\iotest1.txt (10240KiB) |
| 2 | 1904279552 | 464912 | 3.03 | 774.83 | 10.322 | 4.645 | C:\iotest1.txt (10240KiB) |
| 3 | 1446027264 | 353034 | 2.30 | 588.37 | 13.595 | 7.184 | C:\iotest1.txt (10240KiB) |
| ----- | | | | | | | |
| total: | 6698373120 | 1635345 | 10.65 | 2725.50 | 11.739 | 6.076 | |



iotest_10M_512.txt

Diskspd run 2 – 10MB @ 512:

Host Ports:

32048 IOPS
3810 MB/s

Disks:

| Tier I/O Information | | | | | | |
|----------------------|------|------|------|------|-------|-------|
| | A | | B | | Total | |
| Tier | IOPS | MB/s | IOPS | MB/s | %IOPS | %MB/s |
| Performance | 1832 | 1255 | | | 100% | 100% |
| Standard | 0 | 0 | | | 0% | 0% |
| Total | 1832 | 1255 | 0 | 0 | | |

Total Results from within a VM:

Total IO

thread | bytes | I/Os | MiB/s | I/O per s | AvgLat | LatStdDev | file

| | | | | | | | |
|--------|-------------|-------|-------|--------|---------|---------|---------------------------|
| 0 | 9937879040 | 18955 | 15.80 | 31.59 | 253.243 | 103.607 | C:\iotest1.txt (10240KiB) |
| 1 | 9942597632 | 18964 | 15.80 | 31.61 | 253.130 | 99.337 | C:\iotest1.txt (10240KiB) |
| 2 | 9941024768 | 18961 | 15.80 | 31.60 | 253.157 | 100.802 | C:\iotest1.txt (10240KiB) |
| 3 | 9946791936 | 18972 | 15.81 | 31.62 | 253.013 | 102.359 | C:\iotest1.txt (10240KiB) |
| ----- | | | | | | | |
| total: | 39768293376 | 75852 | 63.21 | 126.42 | 253.136 | 101.539 | |



iotest_10M_512.txt

Diskspd run 3 – 2GB @ 512:

Host Ports:

16777 IOPS
1035 MB/s

Disks:

| Tier I/O Information | | | | | | |
|----------------------|--------------|-------------|----------|----------|-------|-------|
| Tier | A | | B | | Total | |
| | IOPS | MB/s | IOPS | MB/s | %IOPS | %MB/s |
| Performance | 12814 | 1154 | | | 100% | 100% |
| Standard | 0 | 0 | | | 0% | 0% |
| Total | 12814 | 1154 | 0 | 0 | | |

Total Results from within a VM:

Total IO

thread | bytes | I/Os | MiB/s | I/O per s | AvgLat | LatStdDev | file

```
-----  
0 | 2709520384 | 5168 | 4.31 | 8.61 | 928.716 | 322.991 | C:\iotest2.txt (2048MiB)  
1 | 2710568960 | 5170 | 4.31 | 8.62 | 928.596 | 326.705 | C:\iotest2.txt (2048MiB)  
2 | 2712666112 | 5174 | 4.31 | 8.62 | 927.955 | 328.458 | C:\iotest2.txt (2048MiB)  
3 | 2711093248 | 5171 | 4.31 | 8.62 | 928.512 | 323.824 | C:\iotest2.txt (2048MiB)  
-----
```

total: 10843848704 | 20683 | 17.24 | 34.47 | 928.445 | 325.503



iotest_2G_512.txt

While running the diskspd test of the second stage, a latency of >200ms was observed. Upon further investigation the latency on the ME4 never exceeded 1ms and the ESXi hosts reported a <5ms latency in “esxstop”. Only within Windows, diskspd and TaskManager, a latency of >200ms is observed. To completely exclude the ME4, the 10 disk Raid 6 was removed and the system re-benchmarked, the latency was still observable, and the Performance-Tier held up on its own perfectly fine. When benchmarking the system through an RDP session to the VDI machines the latency was on average 5ms. The latency of >200ms was only observed while connecting to the VDI machines using the VMware Horizon View Client. Neither the Horizon connection server nor the external host launching the sessions showed resource constrains. As the latency is only observed within Windows and while being connected using VMware Horizon View Client, it leads to the conclusion that the latency is a result of a protocol overhead. A google search resulted in a document describing an optimization case study from VMware, where Windows 7 and 10 are being compared. In this document a latency of 1000ms is being reported for Windows 10. The document can be found under the following link:

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmware-view-horizon-7-windows-10-optimize-virtual-desktop.pdf>

8. Conclusion

As expected, the tiering aspect of the ME4 works flawlessly. Furthermore, it becomes clear that the ME4 is more than capable of handling at least 50 medium VMs / medium users on a single 2 SSD R1 while still having enough performance, even under the heaviest loads. Therefore the minimum configuration using ME4 in conjunction with Horizon VDI setups would be with 2 SSDs, from where the user can upscale it with more SSDs for capacity or additional performance to support their VDI solution or start populating it with capacity disks and start using Pool B for any applications that may be hosted in the environment, while not affecting the performance of the VDI setup.

Something worth mentioning, is that our 2 Hosts were CPU bound, with 50 VDI VMs + infrastructure VMs, we started seeing slight congestions on the Hosts CPU allocations. Which is suggesting that before we need to upscale the storage with more SSDs, we would need to add more CPU power on the hosts.

VMs are always on from time of creation. Uptime of the VM is 100%, provided the user does not sign out of the Windows side, then the VM gets scrapped and re-created from the template, done in typically less then 2min.