



The science behind the report:

Lenovo ThinkAgile CP Series: A fully featured turnkey private cloud, with fast deployment and easy management

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report [“Lenovo ThinkAgile CP Series: A fully featured turnkey private cloud, with fast deployment and easy management”](#).

On August 31, 2018, we finalized the hardware and software configurations we tested. Updates for current and recently released hardware and software appear often, so unavoidably these configurations may not represent the latest versions available when this report appears. We concluded hands-on testing on September 5, 2018.

At the time of this study, we used the Lenovo ThinkAgile CP in conjunction with Cloudistics documentation and software resources. As of H1 2019, Lenovo rebranded portions of the ThinkAgile CP product line. The Cloudistics Marketplace as noted in this report is now called the Lenovo Cloud Marketplace. Our testing reflects the use of the marketplace as it existed in 2018.

Cost analysis for a public cloud platform from Amazon

We calculated costs for Amazon EC2 instances that would match the CPU and memory requirements of VMs that ThinkAgile CP can support. We did not test performance of ThinkAgile CP, so our calculations are theoretical, based on CPU, RAM, and storage capacity of the ThinkAgile CP, assumptions about the workloads that the VMs would support, and decisions about the number of virtual CPUs each physical CPU could support. We compared the AWS cost to the \$492,501 discounted cost of ThinkAgile CP. That price, provided to us by Cloudistics, includes three-year 24x7x4 replacement service, software maintenance for three years, hardware installation, deployment services, and training.

100% three-year, all-up-front reserved cost of Lenovo ThinkAgile CP Series vs. Amazon Web Services™ (AWS)

40-to-60 m5Large:r5large mix	
vCPU to pCPU ratio	3:1
m5.large to r5.large ratio	40:60
Total number of Instances with 2vCPU at 80% capacity	192
AWS	\$602,476
Lenovo platform cost	\$492,501
Savings with Lenovo	\$109,975
Percent savings	18.3%

90% on-demand; remainder 3-year all-upfront reserved cost of Lenovo ThinkAgile CP Series vs Amazon AWS

	3:1	4:1	5:1	6:1	7:1	8:1
vCPU to pCPU ratio	3:1	4:1	5:1	6:1	7:1	8:1
m5.large / r5.large / c5.large mix	40:60:0	80:20:0	30:20:50	0:19:81	0:12:88	0:6:94
vCPU:pCPU	3:1	4:1	5:1	6:1	7:1	8
Total number of Instances with 2vCPU at 80% capacity	192	256	320	384	448	512
AWS	\$834,811	\$925,267	\$1,026,136	\$1,127,579	\$1,227,306	\$1,322,284
Lenovo platform cost	\$492,501	\$492,501	\$492,501	\$492,501	\$492,501	\$492,501
Savings with Lenovo	\$342,310	\$432,766	\$533,635	\$635,078	\$734,805	\$829,783
Percent savings	41.0%	46.8%	52.0%	56.3%	59.9%	62.8%

Usage scenario

We designed a scenario in which an enterprise is planning to run a mix of demanding workloads such as web servers, enterprise applications, and data analytics.

Workloads

We planned for a mix of three different workloads based on instance models available on AWS. The ThinkAgile CP Series runs Intel® Xeon® Scalable processors, so we chose from AWS instances that run on processors from the same family. The instances we selected use two virtual processors each because of the demanding nature of the applications described in the usage scenario—but they have different amounts of memory to accommodate the needs of different applications.

	Model	vCPU	Mem (GiB)
Compute optimized	c5.large	2	4
General purpose	m5.large	2	8
Memory optimized	r5.large	2	16

Available capacity of Lenovo ThinkAgile CP Series at 80% utilization

We sized the VM capacity of the Lenovo ThinkAgile CP Series based on an average of 80 percent utilization of the RAM and physical CPU.

	pCPU	RAM (GB)	Storage (TB)
4-Node Capacity	160	3072	89.6
80% capacity RAM and pCPU	128	2458	N/A

vCPU-to-pCPU ratios

Each VM requires one or more vCPUs. ESXi supports CPU over-committing, which is allocation of more virtual CPUs than the number physical CPUs available on the server. There is no consistent rule of thumb or best practice for a vCPU-to-pCPU ratio. A quick search can yield recommendations ranging from 1:1 to 20:1 or higher. We went with a middle range, modeling solutions based on virtual CPU to physical CPU ratios that range from 3:1 to 8:1.

Identifying a mix of instance types

Here are the steps we used to identify a mix of instance types that came close to an amount of memory and number of vCPUs equal to 80% capacity for RAM and pCPU on ThinkAgile CP Series:

- vCPU at 80% capacity – Our earlier calculation showed 128 pCPU available at 80 percent utilization. We multiplied the vCPU ratio by that value to get the number of vCPUs available at 80 percent capacity..
- Total VMs with two vCPUs – we divided those values by two, the number of vCPU per VM, to get a number of VMs. We priced that number of AWS instances.
- Percentage instance types, number instance types, and RAM usage – our earlier calculation showed 2458 GB RAM at 80 percent utilization. We created a mix of the three instance types that came close to using that memory amount and equaled the Total VMs with 2vCP. This mix is different for each vCPU count. There is enough memory capacity for the lower CPU counts to include all r5.large (16GB) and m5.large (8GB) instances. There is little more than 4 GB available on average for the 8 vCPU-to-1 pCPU instances, so its mix is almost entirely c5.large (4GB instances).

Ratio (vCPU:pCPU)	3	4	5	6	7	8
vCPU at 80% capacity	384	512	640	768	896	1024
Total VMs with 2vCPU at 80% capacity	192	256	320	384	448	512
Percentage instance types						
% m5.large (8GB)	40%	80%	30%	0%	0%	0%
% r5.large (16GB)	60%	20%	20%	19%	12%	6%
% c5.large (4GB)	0%	0%	50%	81%	88%	94%
Number instances types						
m5.large (8GB)	77	205	96	0	0	0
r5.large (16GB)	115	51	64	73	54	31
c5.large (4GB)	0	0	160	311	394	481
RAM usage						
m5.large (8GB)	616	1640	768	0	0	0
r5.large (16GB)	1840	816	1024	1168	864	496
c5.large (4GB)	0	0	640	1244	1576	1924
Total	2456	2456	2432	2412	2440	2420

Calculating costs

We used the AWS Simple Monthly Calculator to get estimates of three-year costs for each of the six vCPU:pCPU ratios we compared. This required us to make decisions about software, payment terms, and other calculator inputs:

- We chose instances with just Linux OS software
- We chose a usage of 80 percent utilized per month.
- We chose US East (Northern Virginia) for the location.
- We made 90-percent instances three-year on-demand and the rest three-year all-upfront reserved. You pay monthly based on a three-year contract for the on-demand; you pay upfront for the all-upfront reserved. The on-demand instances are more expensive but can be commissioned and decommissioned as needed. The reserved instances are best for long-term usage at a steady state. (table below shows the counts of instances based on those percentages. We also show prices for the 3 vCPU:1pCPU ratio using only the least expensive three-year all upfront reserved instances.
- To match the storage on ThinkAgile CP (which uses RAID50 with one hot spare per disk group), we included 90 x 1TB General Purpose SSD (gp2) Amazon EBS volumes totaling 90TB.
- We included business support in the platform.

We used the AWS Simple Monthly Calculator to get estimates of three-year costs for each of the six vCPU-to-pCPU ratios. Because we chose a mix of upfront and monthly billing options, the AWS Simple Monthly Calculator quoted an initial one-time payment and a monthly payment price. We calculated a three-year total price for the configuration we selected along with AWS Business support. We accessed the [AWS Simple Monthly Calculator](#) on 9/21/2018 for the six quotes that include three-year on-demand instances, and on 10/05/2018 for the quote with 00% all upfront reserved instances.

Ratio (vCPU:pCPU)	3	4	5	6	7	8
AWS price options on-demand						
m5.large	69	185	86	0	0	0
r5.large	104	46	58	66	49	28
c5.large	0	0	144	280	355	433
AWS price options up-front reserved						
m5.large	8	20	10	0	0	0
r5.large	11	5	6	7	5	3
c5.large	0	0	16	31	39	48
Estimate of bill						
Total one-time payment	\$23,669.47	\$27,707.65	\$32,701.34	\$36,822.98	\$41,002.40	\$46,053.87
Total monthly payment	\$22,531.72	\$24,932.20	\$27,595.41	\$30,298.78	\$32,952.88	\$35,450.84
3-year total	\$834,811.39	\$925,266.85	\$1,026,136.10	\$1,127,579.06	\$1,227,306.08	\$1,322,284.11

How we tested management tasks within the ThinkAgile CP Series platform

We tested a two-node Lenovo ThinkAgile CP Series for our use cases due to availability. This included 80 vCPU, 256 GB RAM, and 34TB storage. Note that a standard Lenovo ThinkAgile CP Series four-node model would typically have 160 vCPU, 3072 GB RAM, and 57TB Storage. The use cases cover only time and steps spent within the UI; therefore, performance between two-node and four-node is irrelevant in this study.

Results of our hands-on testing

The table below presents time and steps for our work with managing the ThinkAgile CP Series platform. We used a single IT generalist to complete all management tasks within the Lenovo ThinkAgile CP Series platform. The technician received a few hours of training from the Cloudistics engineers who came to our data center. The technician timed each task using a stopwatch. He performed each task three times. We report only the median time for each task.

Use case	Steps	Admin time (seconds)
Simple networking		
Creating an additional Virtual Datacenter with differing resource capacities	3	16.5
Adding a virtual network that has access to the internet	6	51.4
Creating and applying a firewall rule	4	27.2
Application provisioning and lifecycle activities		
Downloading a new Organization Template from the Cloudistics Marketplace	6	21.5
Deploying a Windows VM	6	28.6
Deploying three Windows VMs	6	28.6
Decommissioning a Windows VM	3	11.3
Reallocating VM resources	6	24.1
Allocating a new user	3	14.9
Metering and billing		
Creating a new report	4	58.5
Exporting a report	3	13.2
Simple infrastructure patching		
All patching	0 (Lenovo takes care of this)	0.0 (Lenovo takes care of this)

Simple networking, including multi-tenancy and daily networking maintenance operations

Creating an additional Virtual Datacenter with differing resource capacities

1. From the Cloudistics dashboard, click the + icon next to Infrastructure, and click New Virtual Datacenter from the dropdown.
2. Enter a Virtual Datacenter Name.
3. Review the CPU, Memory, and Storage Resource Summary to determine the resources available for allocation. Enter a number of CPU cores, gigabytes of Memory, and terabytes of Storage to allocate for the VDC, and click Create Virtual Datacenter.

Adding additional networks for virtual machines

Adding a virtual network that has access to the internet

1. From the Cloudistics dashboard, click Networks under Networking.
2. Click New Network, and click New VNET from the dropdown.
3. Enter a VNET Name, Network Address, Netmask, and Default Gateway.
4. Enter a DHCP range for the network.
5. Select the Virtual Datacenter where the new network function virtualization will be deployed. Leave the default 1 CPU, 1GB Memory provisioning. In the Outside Interface (Routing Service) section, select the VNET networking mode, and select the test network.
6. For the Firewall Settings, select the WEB Firewall Profile, and click Create VNET.

Creating and applying a firewall rule

1. From the Cloudistics dashboard, click Firewall Profiles under Networking.
2. Click New Firewall Profile.
3. Enter a Firewall Profile Name.
4. Under the Rules section, leave the default "Allow Incoming TCP" conditions, and leave the From Source IP Range field blank. Enter Source Port Range 443, Destination IP Range 198.168.0.100-198.168.0.150, and Destination Port Range 443. Click Create Firewall Profile.

Application provisioning and lifecycle activities

Downloading a new Organization Template from the Cloudistics Marketplace

1. From the Cloudistics dashboard, click Organization Templates.
2. Click the link to Cloudistics Marketplace.
3. Choose a new Operating System to download. For our test, we used Centos 7.4 (64-bit). Click the template icon.
4. Click the Download Template button to confirm the selection.
5. From the Download Template menu, select a Save Location. For our test, we selected the Main_VDC template location.
6. Leave the default Template Name, Description, CPU and Memory Provisioning, and Boot Order, and click Save Application Template. The Organization Template will begin downloading and be added to the inventory.

Deploying a Windows VM

1. From the Cloudistics dashboard, under Applications, click Organization Templates.
2. From the template menu, next to the Windows 10 Enterprise (64-bit) template, click New instance.
3. From the deployment screen, configure the virtual machine options. Enter an Instance Name.
4. Choose a Virtual Datacenter for deployment. Optionally choose an Application Group, and number of instances for deployment.
5. Leave the default Storage Pool and Migration Zone. Select VNET for Networking Mode, and choose a Network from the dropdown. Leave the Default Firewall Override and Mac Address options.
6. Leave the default Instance Settings, and click Create Instance.

Deploying three Windows VMs

1. From the Cloudistics dashboard, under Applications, click Organization Templates.
2. From the template menu, next to the Windows 10 Enterprise (64-bit) template, click New instance.
3. From the deployment screen, configure the virtual machine options. Enter an Instance Name.
4. Choose a Virtual Datacenter for deployment. Optionally choose an Application Group, and number of instances for deployment.
5. Leave the default Storage Pool and Migration Zone. Select VNET for Networking Mode, and choose a Network from the dropdown. Leave the Default Firewall Override and Mac Address options.
6. Leave the default Instance Settings, and click Create Instance.

Decommissioning a Windows VM

1. From the Cloudistics dashboard, under Applications, click All Instances.
2. Ensure the virtual machine is turned off. With the virtual machine turned off, click to highlight the virtual machine to be decommissioned.
3. Click the "... " icon, and click Delete from the dropdown. Click Delete Application Instances to confirm decommissioning. The virtual machine will be deleted.

Reallocating VM resources

1. From the Cloudistics dashboard, under Applications, click All Instances.
2. Ensure the virtual machine is shutdown. Click the "... " dropdown menu icon, hover over Edit Instance, and click Edit CPU Provisioning.
3. Change the CPU core count to 4, and click Save.
4. Click the "... " dropdown menu icon again, hover over Edit Instance, and click Edit Memory Provisioning.
5. Change the memory GiB count to 4, and click Save.
6. Select the Instance by clicking on its field, and click Start/resume to restart the virtual machine. Click Start/Resume Instances to confirm the action.

Allocating a new user

1. From the Cloudistics Dashboard, click User Management under General.
2. Click Invite User.
3. Enter an Email Address, and chose an Access Level. For our testing, we chose Infrastructure Viewer. Click Send invitation.

Metering and billing tasks

Creating a new report

1. From the Cloudistics Metering Manager dashboard, click + Add next to report.
2. Enter a report Name, choose a Scheduled Interval, choose a Report type, and choose an Output type. For our testing, we chose an Hourly interval, Detail report type, and Local Database Only.
3. Add a meter for each metric you wish to track. For our testing, we added a VCPU, RAM, and Storage Meter. For each Meter, enter a name, choose a metric, measurement (allocation or utilization), enter a rate, and rate interval.
4. Choose the applications to run the report on, and alternatively choose an entire Data Center. For our testing, we chose all deployed applications in our datacenter. Click Save to complete the report.

Exporting a report

1. From the Cloudistics Metering Manager dashboard, click Reports.
2. Click the Name of the test report.
3. Scroll to the Report Runs section. Click the XLS icon of the most recent report run to download the report as an Excel document.

Read the report at <http://facts.pt/8j2vwr> ►

This project was commissioned by Lenovo.



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners.

DISCLAIMER OF WARRANTIES; LIMITATION OF LIABILITY:

Principled Technologies, Inc. has made reasonable efforts to ensure the accuracy and validity of its testing, however, Principled Technologies, Inc. specifically disclaims any warranty, expressed or implied, relating to the test results and analysis, their accuracy, completeness or quality, including any implied warranty of fitness for any particular purpose. All persons or entities relying on the results of any testing do so at their own risk, and agree that Principled Technologies, Inc., its employees and its subcontractors shall have no liability whatsoever from any claim of loss or damage on account of any alleged error or defect in any testing procedure or result.

In no event shall Principled Technologies, Inc. be liable for indirect, special, incidental, or consequential damages in connection with its testing, even if advised of the possibility of such damages. In no event shall Principled Technologies, Inc.'s liability, including for direct damages, exceed the amounts paid in connection with Principled Technologies, Inc.'s testing. Customer's sole and exclusive remedies are as set forth herein.