

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

CIRBA INC. (d/b/a DENSIFY)
and CIRBA IP, INC.,

Plaintiffs/Counter-Defendants,

v.

VMWARE, INC.,

Defendant/Counter-Plaintiff.

C.A. No. 19-742-LPS

(Consolidated)

JOINT CLAIM CONSTRUCTION BRIEF

Dated: November 22, 2021

OF COUNSEL:

Arturo J. González
Michael A. Jacobs
Richard S.J. Hung
MORRISON & FOERSTER LLP
425 Market Street
San Francisco, CA 94105
(415) 268-7000
agonzalez@mofo.com
mjacobs@mofo.com
rhung@mofo.com

Bitra Rahebi
MORRISON & FOERSTER LLP
707 Wilshire Boulevard
Los Angeles, CA 90017
(213) 892-5200
brahebi@mofo.com

Scott F. Llewellyn
MORRISON & FOERSTER LLP
4200 Republic Plaza
370 Seventeenth Street
Denver, CO 80202-5638
(303) 592-2204
sllewellyn@mofo.com

YOUNG CONAWAY STARGATT &
TAYLOR, LLP

Anne Shea Gaza (No. 4093)
Robert M. Vrana (No. 5666)
Samantha G. Wilson (No. 5816)
Rodney Square
1000 North King Street
Wilmington, DE 19801
(302) 571-6600
agaza@ycst.com
rvrana@ycst.com
swilson@ycst.com

Attorneys for VMware, Inc.

Sarah O. Jorgensen (*pro hac vice*)
sjorgensen@reichmanjorgensen.com
REICHMAN JORGENSEN LEHMAN &
FELDBERG LLP
1201 West Peachtree, Suite 2300
Atlanta, GA 30309
Telephone: (404) 609-1040
Telecopier: (650) 623-1449

Christine E. Lehman (*pro hac vice*)
clehman@reichmanjorgensen.com
REICHMAN JORGENSEN LEHMAN &
FELDBERG LLP
1710 Rhode Island Avenue, 12th Floor
Washington, DC 20036
Telephone: (202) 894-7310
Telecopier: (650) 623-1449

Khue V. Hoang (*pro hac vice*)
khoang@reichmanjorgensen.com
REICHMAN JORGENSEN LEHMAN &
FELDBERG LLP
750 Third Avenue, Suite 2400
New York, NY 10017
Telephone: (212) 381-1965
Telecopier: (650) 623-1449

Kenneth L. Dorsney (#3726)
kdorsney@morrisjames.com
Cortlan Hitch (#6720)
chitch@morrisjames.com
MORRIS JAMES LLP
500 Delaware Avenue, Suite 1500
Wilmington, DE 19801
Telephone: (302) 888-6800

Courtland L. Reichman (*pro hac vice*)
creichman@reichmanjorgensen.com
REICHMAN JORGENSEN LEHMAN &
FELDBERG LLP
100 Marine Parkway, Suite 300
Redwood Shores, California 94065
Telephone: (650) 623-1401
Telecopier: (650) 623-1449

*Attorneys for Cirba Inc.
(d/b/a Densify) and Cirba IP, Inc.*

VMWARE’S TABLE OF ABBREVIATIONS

Term	Abbreviation
U.S. Patent No. 8,209,687	’687 patent
U.S. Patent No. 10,523,492	’492 patent
U.S. Patent No. 10,951,459	’459 patent
Distributed Resource Scheduler	DRS
Virtual Machine or Virtual Guest	VM
Virtual Host	host
Patent Trial and Appeal Board	PTAB
<i>Inter Partes</i> Review	IPR
computer-readable storage medium comprising instructions	CRM
United States Patent and Trademark Office	USPTO or PTO
person of ordinary skill in the art	POSITA
Trial Transcript (D.I. 587 through D.I. 597)	TT

TABLE OF CONTENTS

	Page
I. Agreed Upon Constructions.....	1
A. U.S. Patent No. 10,523,492 (the “492 patent”)	1
B. U.S. Patent No. 10,951,459 (the “459 patent”)	1
II. Disputed Constructions for Densify’s Patents	1
A. Densify’s Level of Ordinary Skill In The Art.....	1
B. VMware’s Level of Ordinary Skill In The Art	1
C. Terms for U.S. Patent No. 8,209,687 (“687 Patent”)	2
1. “evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to said technical, business and workload constraints” (‘687 patent: claims 3, 7)	2
a. Densify’s Opening Position	2
(i) VMware’s Proposed Construction Changes the Meaning of the Claim.	2
(1) The Plain Language Forbids VMware’s Construction.	3
(2) The Specification Doesn’t Require Rule Sets for Each VM.	5
(ii) This Term Has Already Been Fully Construed and Is Law of the Case.....	6
b. VMware’s Answering Position.....	11
(i) The Intrinsic Evidence Supports VMware’s Construction.....	13
(1) The Claim Language Supports VMware’s Construction.....	13
(2) The Specification Supports VMware’s Construction.....	16
(ii) Cirba’s Textual Counterarguments Lack Merit	18
(iii) Cirba’s Procedural Arguments Also Lack Merit	21
c. Densify’s Reply Position	22
(i) VMware’s “Do-Over” Construction is Contrary to the Intrinsic Evidence.	23
(ii) The Plain Language Contradicts VMware’s “Do- Over” Construction.	24

TABLE OF CONTENTS

(continued)

	Page
(1) The Specification Demonstrates The Error in VMware's "Do-Over" Construction.....	26
(2) VMware's Criticisms Of Densify's "Textual Counterarguments" Fail Because They Are Not Supported By The Plain Language.	28
(iii) VMware Fails To Offer Good Cause For Its Changed Construction.....	29
d. VMware's Sur-Reply Position.....	31
(i) Each Evaluation Must Use Rules for Each Constraint.....	31
(ii) VMware's Construction Does Not Require Rules on Each VM.....	32
(iii) VMware's Construction Does Not Require that All Rule Sets Have Three Constraints	32
(iv) All Three Constraints Are Associated with VMs and Are Not Optional.....	32
(v) Cirba's Distinction between "One or More Rules" and "One or More Rule Sets" Is a Distinction without a Difference	33
(vi) VMware Did Not Delay Raising the Need for Claim Construction.....	33
(vii) The Court Did Not Construe the Term Before and Should Do So Now	34
D. Terms Common to U.S. Patent No. 10,951,459 ("459 Patent") and U.S. Patent No. 10,523,492 ("492 Patent")	34
1. "source system" / "source computer system" ('492 patent: claims 1, 12, 23) ('459 patent: claims 1-3, 9-10, 16-17, 22-23, 26-34, 40-41, 47-48, 53-54, 57-63)	34
a. Densify's Opening Position	34
b. VMware's Answering Position.....	35
(i) The Inclusion of "Are Moved" Is Improper	36
(ii) Cirba's Proposed Limitation to "Physical, Virtual, or Hypothetical Systems" Should Be Rejected.....	37
c. Densify's Reply Position	38
(i) The Specification Supports "Are Moved."	39

TABLE OF CONTENTS
(continued)

	Page
(ii) The Patents Disclose Physical, Virtual, and Hypothetical Systems.....	39
d. VMware’s Sur-Reply Position.....	40
(i) Including “Are Moved” Would Be Improper	40
(ii) Limitation to “Physical, Virtual, or Hypothetical” Should Be Rejected.....	42
2. “target system” / “target computer system” (’492 patent: claims 1, 4-6, 9, 11-12, 15-17, 20, 22-23, 26-28, 31, 33) (’459 patent: claims 1-3, 9-10, 16-17, 22-23, 26-34, 40-41, 47-48, 53-54, 57-63)	42
a. Densify’s Opening Position	42
b. VMware’s Answering Position.....	43
c. Densify’s Reply Position	44
d. VMware’s Sur-Reply Position.....	44
3. “ <i>place</i> ” terms (’492 patent: claims 1, 12, 23) (’459 patent: claims 1, 32, 63)	44
a. Densify’s Opening Position	45
b. VMware’s Answering Position.....	47
c. Densify’s Reply Position	50
d. VMware’s Sur-Reply Position.....	52
E. Terms for U.S. Patent No. 10,523,492 (“’492 Patent”) Only	53
1. “the systems” / “the source systems” / “the target systems” (’492 patent: claims 1, 12, 23).....	53
a. Densify’s Opening Position	54
b. VMware’s Answering Position.....	55
(i) “the systems”	55
(ii) “the source systems”	57
(iii) “the target systems”	57
c. Densify’s Reply Position	58
(i) “the systems”	58
(ii) “the source systems”	60
(iii) “the target systems”	60
d. VMware’s Sur-Reply Position.....	61

TABLE OF CONTENTS
(continued)

	Page
(i) “the systems”	61
(ii) “the source systems”	62
(iii) “the target systems”	62
2. “at least one of a compatibility score, and a number of transfers” ('492 patent: claims 3, 14, 25)	63
a. Densify’s Opening Position	63
b. VMware’s Answering Position.....	64
c. Densify’s Reply Position	65
d. VMware’s Sur-Reply Position.....	66
3. “determining at least one optimal combination of applications operable independently on the target system” ('492 patent: claims 9, 20, 31)	67
a. Densify’s Opening Position	67
b. VMware’s Answering Position.....	68
c. Densify’s Reply Position	69
d. VMware’s Sur-Reply Position.....	70
4. “a 1-to-1 compatibility analysis” ('492 patent: claims 1, 12, 23).....	71
a. Densify’s Opening Position	71
b. VMware’s Answering Position.....	72
c. Densify’s Reply Position	74
d. VMware’s Sur-Reply Position.....	74
5. “an N-to-1 compatibility analysis” ('492 patent: claims 1, 12, 23).....	74
a. Densify’s Opening Position	75
b. VMware’s Answering Position.....	76
c. Densify’s Reply Position	78
d. VMware’s Sur-Reply Position.....	79
F. Terms for U.S. Patent No. 10,951,459 (“’459 Patent”) Only	80
1. “the source systems” ('459 patent: claims 1, 31, 63).....	80
a. Densify’s Opening Position	80
b. VMware’s Answering Position.....	82
c. Densify’s Reply Position	83

TABLE OF CONTENTS
(continued)

	Page
d. VMware’s Sur-Reply Position.....	84
III. Disputed Constructions for VMware’s Patents.....	84
A. U.S. Patent No. 9,379,995 (“’995 Patent”).....	84
1. “snapshot” (’995 Patent: claims 1, 9, and 17).....	84
a. VMware’s Opening Position.....	84
b. Densify’s Answering Position	85
c. VMware’s Reply Position.....	88
d. Densify’s Sur-Reply Position	89
2. “iteratively traversing a resource hierarchy” (’995 Patent: claims 1, 9, and 17)	89
a. VMware’s Opening Position.....	89
b. Densify’s Answering Position	92
c. VMware’s Reply Position.....	97
d. Densify’s Sur-Reply Position	100
3. “target resource allocation” (’995 Patent: claims 1, 9, and 17)	101
a. VMware’s Opening Position.....	101
b. Densify’s Answering Position	102
c. VMware’s Reply Position.....	102
d. Densify’s Sur-Reply Position	102
4. “resource configuration action” (’995 Patent: claims 1, 9, and 17).....	103
a. VMware’s Opening Position.....	103
b. Densify’s Answering Position	103
c. VMware’s Reply Position.....	104
d. Densify’s Sur-Reply Position	105
5. “capacity expansion action” (’995 Patent: claims 1, 9, and 17)	105
a. VMware’s Opening Position.....	105
b. Densify’s Answering Position	106
c. VMware’s Reply Position.....	106
d. Densify’s Sur-Reply Position	106
B. U.S. Patent No. 9,766,945 (“’945 Patent”).....	107

TABLE OF CONTENTS
(continued)

	Page
1. “calculating a target resource configuration” (’945 Patent: claims 1, 8, and 15)	107
a. VMware’s Opening Position.....	107
b. Densify’s Answering Position	109
c. VMware’s Reply Position.....	110
d. Densify’s Sur-Reply Position	110
C. U.S. Patent No. 10,025,638 (“’638 Patent”)	110
1. “operationally distinct cloud computing facilities” (’638 Patent: claims 10 and 19).....	110
a. VMware’s Opening Position.....	110
b. Densify’s Answering Position	113
c. VMware’s Reply Position.....	117
d. Densify’s Sur-Reply Position	119
2. “cloud-connector server” (’638 Patent: claims 10 and 19).....	120
a. VMware’s Opening Position.....	120
b. Densify’s Answering Position	122
c. VMware’s Reply Position.....	123
d. Densify’s Sur-Reply Position	124
3. “cloud-connector node” (’638 Patent: claims 10 and 19).....	125
a. VMware’s Opening Position.....	125
b. Densify’s Answering Position	127
c. VMware’s Reply Position.....	128
d. Densify’s Sur-Reply Position	129
D. U.S. Patent No. 10,261,842 (“’842 Patent”)	130
1. “distributed resource scheduling module” (’842 Patent: claim 1)	130
a. VMware’s Opening Position.....	131
b. Densify’s Answering Position	133
(i) Functions Carried Out By The Claimed “Module” Include “Receiving..., Determining..., Calculating..., and Performing...”	134
(ii) The Specification Does Not Disclose Corresponding Structure.	135

TABLE OF CONTENTS
(continued)

		Page
	c. VMware’s Reply Position.....	137
	(i) Distributed Resource Scheduling Modules are Structures Known in the Art and Thus Not a Means Plus Function Claim Term.....	137
	(ii) Even if “Distributed Resource Scheduling Module” Were Subject to § 112, ¶ 6, the ’842 Patent Discloses Sufficient Structure.....	138
	d. Densify’s Sur-Reply Position	142
2.	“A non-transitory computer-readable storage medium comprising instructions that, when executed in a computing device, causes the computing device to carry out the steps of” (claim 12) (’842 Patent: claim 12)	144
	a. VMware’s Opening Position.....	144
	b. Densify’s Answering Position	145
	(i) The Preamble Is Limiting	146
	(ii) Claim 12 Is Subject to 35 U.S.C. § 112, ¶ 6	146
	(iii) The Functions Performed By the “Non-Transitory Computer-Readable Storage Medium Comprising Instructions...” Term Are the Steps Listed In the Claim.....	147
	(iv) The Specification Does Not Disclose Corresponding Structure.	148
	c. VMware’s Reply Position.....	149
	(i) The Claimed Instructions Are the Claimed Structure....	149
	(ii) Even if the CRM Term Were Subject to § 112, ¶ 6, the ’842 Patent Discloses Sufficient Structure.....	151
	d. Densify’s Sur-Reply Position	151
3.	“a virtual management computer configured to” (’842 Patent: claim 17)	152
	a. VMware’s Opening Position.....	153
	b. Densify’s Answering Position	153
	c. VMware’s Reply Position.....	154
	d. Densify’s Sur-Reply Position	154

TABLE OF CONTENTS
(continued)

		Page
4.	“a combined impact on current and future workload” (’842 Patent: claims 1, 12 and 17).....	154
a.	VMware’s Opening Position.....	155
b.	Densify’s Answering Position	156
c.	VMware’s Reply Position.....	157
d.	Densify’s Sur-Reply Position	158

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>3G Licensing, S.A. v. Blackberry Ltd.</i> , No. 17-82-LPS-CJB, 2018 WL 4375091 (D. Del. Sept. 13, 2018)	118
<i>Aatrix Software, Inc. v. Green Shades Software, Inc.</i> , 890 F.3d 1354 (Fed. Cir. 2018).....	9
<i>Akamai Techs., Inc. v. Limelight Nets., Inc.</i> , 805 F.3d 1368 (Fed. Cir. 2015).....	9
<i>Align Tech., Inc. v. 3Shape</i> , No. 17-cv-1648-LPS, 2021 WL 2320139 (D. Del. June 7, 2021)	139
<i>All Dental Prodx, LLC v. Advantage Dental Prod., Inc.</i> , 309 F.3d 774 (Fed. Cir. 2002).....	67
<i>In re Altenpohl</i> , 500 F.2d 1151 (C.C.P.A. 1974)	55
<i>Anderson v. Int’l Eng’g & Mfg., Inc.</i> , 160 F.3d 1345 (Fed. Cir. 1998).....	96
<i>Andover Healthcare, Inc. v. 3M Co.</i> , No. CV 13-843-LPS, 2015 WL 2227786 (D. Del. May 11, 2015).....	86
<i>Apple Inc. v. Motorola, Inc.</i> , 757 F.3d 1286 (Fed. Cir. 2014).....	151
<i>Arendi S.A.R.L. v. LG Elecs., Inc.</i> , 2019 WL 3891150 (D. Del. Aug. 8, 2019) (J. Stark)	<i>passim</i>
<i>Aristocrat Techs. Australia Pty Ltd. v. Int’l Game Tech.</i> , 521 F.3d 1328 (Fed. Cir. 2008).....	143, 151, 152, 154
<i>Arizona v. California</i> , 460 U.S. 605 (1983).....	7, 11
<i>Bancorp Servs., L.L.C. v. Hartford Life Ins. Co.</i> , 359 F.3d 1367 (Fed. Cir. 2004).....	67
<i>Bell & Howell Document Mgmt. Prod. Co. v. Altek Sys.</i> , 132 F.3d 701 (Fed. Cir. 1997).....	86, 87, 88

<i>Bio-Rad Labs., Inc. v. Int’l Trade Comm’n</i> , 998 F.3d 1320 (Fed. Cir. 2021).....	4
<i>Bristol-Myers Squibb Co. v. Aurobindo Pharma USA Inc.</i> , No. CV 17-374-LPS, 2018 WL 5077895 (D. Del. Oct. 18, 2018)	86
<i>CCS Fitness, Inc. v. Brunswick Corp.</i> , 288 F.3d 1359 (Fed. Cir. 2002).....	122, 125
<i>ChefAm., Inc. v. Lamb-Weston, Inc.</i> , 358 F.3d 1371 (Fed. Cir. 2004).....	55
<i>Chevron U.S.A. Inc. v. Univ. of Wyoming Rsch. Corp.</i> , 978 F.3d 1361 (Fed. Cir. 2020).....	85
<i>Cirba Inc. v. VMware, Inc.</i> , No. CV 19-742-LPS, 2020 WL 5910112 (D. Del. Oct. 6, 2020) (J. Stark).....	136
<i>Citrix Sys., Inc. v. Workspot, Inc.</i> , No. 18-588- 2020 WL 5634219 (D. Del. Sept. 21, 2020)	132
<i>Collabo Innovations, Inc. v. Omnivision Techs., Inc.</i> , No. 16-197-JFB-SRF, 2017 WL 3670661 (D. Del. Aug. 25, 2017).....	55
<i>Cordis Corp. v. Boston Sci. Corp.</i> , 561 F.3d 1319 (Fed. Cir. 2009).....	9
<i>Curtiss-Wright Flow Control Corp. v. Velan, Inc.</i> , 438 F.3d 1374 (Fed. Cir. 2006).....	96
<i>Digital–Vending Servs. Int’l, LLC v. Univ. of Phoenix, Inc.</i> , 672 F.3d 1270 (Fed. Cir. 2012).....	54, 80
<i>Douglas Dynamics, LLC v. Buyers Prods. Co.</i> , 717 F.3d 1336 (Fed. Cir. 2013).....	85, 88, 89
<i>Edwards Lifesciences LLC v. Cook Inc.</i> , 582 F.3d 1322 (Fed. Cir. 2009).....	96
<i>Energizer Holdings, Inc. v. Int’l Trade Comm’n</i> , 435 F.3d 1366 (Fed. Cir. 2006).....	54, 55
<i>ePlus, Inc. v. Lawson Software, Inc.</i> , No. 3:09-cv-620, 2010 WL 1779973 (E.D. Va. Apr. 30, 2010)	77
<i>Finisar Corp. v. DirecTV Grp., Inc.</i> , 523 F.3d 1323 (Fed. Cir. 2008).....	143, 152

<i>Finjan, Inc. v. Secure Computing Corp.</i> , 626 F.3d 1197 (Fed. Cir. 2010).....	9
<i>Fonar Corp. v. Johnson & Johnson</i> , 821 F.2d 627 (Fed. Cir. 1987).....	86, 89
<i>Fresenius Med. Care Holding, Inc. v. Baxter Int'l, Inc.</i> , No. C 03-1431 SBA, 2006 WL 1343648 (N.D. Cal. May 16, 2006)	146
<i>Function Media, L.L.C. v. Google, Inc.</i> , 708 F.3d 1310 (Fed. Cir. 2013).....	9, 143
<i>Gindes v. United States</i> , 740 F.2d 947 (Fed. Cir. 1984).....	11
<i>Halliburton Energy Servs., Inc. v. M-I LLC</i> , 514 F.3d 1244 (Fed. Cir. 2008).....	96
<i>Helmsderfer v. Bobrick Washroom Equip., Inc.</i> , 527 F.3d 1379 (Fed. Cir. 2008).....	66
<i>Hill-Rom Servs., v. Stryker Corp.</i> , 755 F.3d 1367 (Fed. Cir. 2014).....	98, 128
<i>In re Hinker Co.</i> , 150 F.3d 1362 (Fed. Cir. 1998).....	2
<i>Honeywell Inc. v. Victor Co. of Japan</i> , 298 F.3d 1317 (Fed. Cir. 2002).....	122
<i>Housey Pharms., Inc. v. Astrazeneca UK Ltd.</i> , 366 F.3d 1348 (Fed. Cir. 2004).....	85
<i>HTC Corp. v. IPCom GmbH & Co., KG</i> , 667 F.3d 1270 (Fed. Cir. 2012).....	14
<i>HZNP Medicines LLC v. Actavis Lab'ys UT, Inc.</i> , 940 F.3d 680 (Fed. Cir. 2019).....	95, 116
<i>Infinity Computer Prods., Inc. v. Oki Data Americas, Inc.</i> , 987 F.3d 1053 (Fed. Cir. 2021).....	157, 158
<i>IpLearn, LLC v. Kenexa Corp.</i> , No. 11-825-RGA, 2013 WL 5730610 (D. Del. Oct. 22, 2013)	85, 88
<i>IPXL Holdings, LLC v. Amazon.com, Inc.</i> , 333 F. Supp. 2d 513 (E.D. Va. 2004)	66

<i>Irdeto Access, Inc. v. Echostar Satellite Corp.</i> , 383 F.3d 1295 (Fed. Cir. 2004).....	104
<i>Jack Guttman, Inc. v. Kopykake Enterprises, Inc.</i> , 302 F.3d 1352 (Fed. Cir. 2002).....	122
<i>Kaneka Corp. v. Xiamen Kingdomway Grp. Co.</i> , 790 F.3d 1298 (Fed. Cir. 2015).....	6
<i>Kraft Foods, Inc. v. Int’l Trading Co.</i> , 203 F.3d 1362 (Fed.Cir.2000).....	87
<i>Kruse Tech. P’ship v. Volkswagen AG</i> , 544 Fed. App’x 943 (Fed. Cir. 2013).....	11
<i>Liebel-Flarsheim Co. v. Medrad, Inc.</i> , 358 F.3d 898 (Fed. Cir. 2004).....	45
<i>Liebscher v. Boothroyd</i> , 258 F.2d 948 (C.C.P.A. 1958)	48
<i>LifeNet Health v. LifeCell Corp.</i> , 837 F.3d 1316 (Fed. Cir. 2016).....	9
<i>Linear Tech. Corp. v. Impala Linear Corp.</i> , 379 F.3d 1311 (Fed. Cir. 2004).....	137
<i>Max’s Seafood Café ex rel. Lou-Ann, Inc. v. Quinteros</i> , 176 F.3d 669 (3d Cir. 1999).....	7, 29
<i>M2M Sols. LLC v. Sierra Wireless Am., Inc.</i> , No. 12-30-RGA, 2015 WL 5826816 (D. Del. Oct. 2, 2015)	131, 144
<i>M2M Sols. LLC v. Sierra Wireless Am., Inc.</i> , No. 12-30-RGA, 2016 WL 1298961 (D. Del. Mar. 31, 2016)	<i>passim</i>
<i>Magna Elecs., Inc. v. TRW Auto. Holdings Corp.</i> , No. 12-cv-654, 2015 WL 11401855 (W.D. Mich. Apr. 28, 2015)	117
<i>Markman v. Westview Instruments, Inc.</i> , 517 U.S. 370 (1996).....	8, 96
<i>Martek Biosciences Corp. v. Nutrinova, Inc.</i> , 579 F.3d 1363 (Fed. Cir. 2009).....	85, 89, 128
<i>Media Rts. Techs., Inc. v. Cap. One Fin. Corp.</i> , 800 F.3d 1366 (Fed. Cir. 2015).....	92

<i>Nautilus, Inc. v. Biosig Instruments, Inc.</i> , 572 U.S. 898, 134 S. Ct. 2120 (2014).....	54, 69, 81, 92, 96, 116, 159
<i>Noah Sys., Inc. v. Intuit Inc.</i> , 675 F.3d 1302 (Fed. Cir. 2012).....	134
<i>Northern Telecom Ltd. v. Samsung Elecs. Co., Ltd.</i> , 215 F.3d 1281 (Fed. Cir. 2000).....	4
<i>Nuance Commc 'ns, Inc. v. ABBYY USA Software House, Inc.</i> , 813 F.3d 1368 (Fed. Cir. 2016).....	7, 8, 9, 11, 29
<i>O2 Micro International v. Beyond Innovation Technology Company</i> , 521 F.3d 1351 (Fed. Cir. 2008).....	7, 8, 12
<i>Omega Eng'g, Inc. v. Raytek Corp.</i> , 334 F.3d 1314 (Fed. Cir. 2003).....	124
<i>In re Packard</i> , 751 F.3d 1307 (Fed. Cir. May 6, 2014)	55
<i>Personalized Media Commc 'ns, LLC v. ITC</i> , 161 F.3d 696 (Fed. Cir. 1998).....	137, 142
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005) (en banc).....	47, 67, 80, 96
<i>Pitney Bowes, Inc. v. Hewlett-Packard Co.</i> , 182 F.3d 1298 (Fed. Cir. 1999).....	38
<i>Retractable Techs., Inc. v. Becton, Dickinson & Co.</i> , 653 F.3d 1296 (Fed. Cir. 2011).....	97
<i>Rowe v. Dror</i> , 112 F.3d 473 (Fed. Cir. 1997).....	146
<i>S3 Inc. v. NVIDIA Corp.</i> , 259 F.3d 1364 (Fed. Cir. 2001).....	118
<i>Shire Dev. LLC v. Teva Pharms. USA, Inc.</i> , No. 1:17-CV-01696-RGA, 2019 WL 969638 (D. Del. Feb. 28, 2019).....	123, 130
<i>SIMO Holdings, Inc. v. Hong Kong uCloudlink Network Tech. Ltd.</i> , 983 F.3d 1367 (Fed. Cir. 2021).....	64, 65, 66
<i>In re Skvorecz</i> , 580 F.3d 1262 (Fed. Cir. 2009).....	118

<i>Slimfold Mfg. Co. v. Kinkead Indus., Inc.</i> , 810 F.2d 1113 (Fed. Cir. 1987).....	55
<i>Sonix Tech. Co. v. Publ'ns Int'l, Ltd.</i> , 844 F.3d 1370 (Fed. Cir. 2017).....	119
<i>Springs Window Fashions LP v. Novo Indus., L.P.</i> , 323 F.3d 989 (Fed. Cir. 2003).....	127
<i>SuperGuide Corp. v. DirecTV Enters., Inc.</i> , 358 F.3d 870 (Fed. Cir. 2004).....	64, 66
<i>Symantec Corp. v. Computer Assocs. Int'l, Inc.</i> , 522 F.3d 1279 (Fed. Cir. 2008).....	145
<i>Teva Pharms. USA, Inc. v. Sandoz, Inc.</i> , 574 U.S. 318 (2015).....	8
<i>TQ Delta, LLC v. 2Wire, Inc.</i> , 2018 WL 626472 (D. Del. Jan. 30, 2018).....	134
<i>Trading Techs. Int'l, Inc. v. eSpeed, Inc.</i> , 595 F.3d 1340 (Fed. Cir. 2010).....	38
<i>Trimble Inc. v. PerDiemCo LLC</i> , 997 F.3d 1147 (Fed. Cir. 2021).....	9
<i>Uniloc 2017 LLC v. Google LLC</i> , No. 2:18-cv-492-JRG-RSP, 2020 WL 569858 (E.D. Tex. Feb. 5, 2020).....	150, 151
<i>Vanderlande Indus. Nederland BV v. Int'l Trade Comm'n</i> , 366 F.3d 1311 (Fed. Cir. 2004).....	38
<i>Verizon Servs. Corp. v. Cox Fibernet Va., Inc.</i> , 602 F.3d 1325 (Fed. Cir. 2010).....	9
<i>Vivus, Inc. v. Actavis Lab'ys FL Inc.</i> , No. 14-3786, 2016 WL 3919455 (D.N.J. July 20, 2016)	73
<i>VMware, Inc., v. Cirba IP Inc.</i> , IPR2021-00008, Paper 8 (P.T.A.B. Feb. 18, 2021).....	38
<i>Waldorf v. Shuta</i> , 142 F.3d 601 (3d Cir. 1998).....	7
<i>Williamson v. Citrix Online, LLC</i> , 792 F.3d 1339 (Fed. Cir. 2015).....	<i>passim</i>

<i>World Class Tech. Corp. v. Ormco Corp.</i> , 769 F.3d 1120 (Fed. Cir. 2014).....	86
<i>Zeroclick, LLC v. Apple Inc.</i> , 891 F.3d 1003 (Fed. Cir. 2018).....	131, 133, 144

Statutes

35 U.S.C. § 101	<i>passim</i>
-----------------------	---------------

Other Authorities

D. Kappos, “Subject Matter Eligibility of Computer Readable Media,” 1351 Off. Gaz. Pat. Office 212 (Feb. 23, 2010).....	145
Fed. R. Civ. P. 16(b)(4).....	8
Fed. R. Civ. P. 60(b)	8, 9
M.P.E.P § 2181	133
M.P.E.P § 2111.05. Sec. III	144
M.P.E.P § 2173.05(e).....	55

I. AGREED UPON CONSTRUCTIONS**A. U.S. Patent No. 10,523,492 (the “’492 patent”)**

Claim Term/Phrase	Agreed-Upon Proposed Construction
<u>Claims 1, 12, 23</u> “an N-by-N compatibility analysis”	“an analysis that evaluates the compatibility of each of N source systems against each of the other source systems among the N source systems”

B. U.S. Patent No. 10,951,459 (the “’459 patent”)

Claim Term/Phrase	Agreed-Upon Proposed Construction
<u>Claims 1, 31, 63</u> “the source and target systems being evaluated”	“the specific source system and the specific target system being evaluated”
<u>Claims 1, 63</u> “the specific any onc [sic] of the plurality of target system”	“the specific target system”

II. DISPUTED CONSTRUCTIONS FOR DENSIFY’S PATENTS**A. Densify’s Level of Ordinary Skill In The Art**

Densify submits that the appropriate level for one of ordinary skill in the art as of the date of the invention of the respective patents would be an individual with at least a bachelor’s degree in Computer Engineering, Computer Science or comparable degree, and two years of experience developing software tools and/or computer systems for use in the area of information technology infrastructures and utilities in designing and evaluating virtualized environments.

B. VMware’s Level of Ordinary Skill In The Art

For claim construction purposes, there is no material difference between VMware’s level of ordinary skill in the art for the ’687 patent (D.I. 270 at 6) and Cirba’s. For the ’492 and ’459 patents, however, VMware disagrees that a person of ordinary skill in the art (“POSITA”) would

have experience in designing and evaluating virtualized environments, which are not the focus of those patents. Instead, the person of ordinary skill in the art would have experience in designing and evaluating physical (non-virtualized) computing environments. Presently, VMware does not believe that this difference is relevant to construing the disputed terms.

C. Terms for U.S. Patent No. 8,209,687 (“’687 Patent”)

1. “evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to said technical, business and workload constraints” (’687 patent: claims 3, 7)

VMware’s Proposed Construction	Densify’s Proposed Construction
“evaluating each virtual machine against each virtual host and each virtual machine against other virtual machines, in each case using one or more rules pertaining to each of technical, business and workload constraints associated with the corresponding virtual machine”	<p>This limitation has already been fully construed by the Court (see D.I. 356) as:</p> <p>“evaluating each virtual machine against each virtual host and other virtual machines using one or more rule sets pertaining to said technical, business and workload constraints”</p> <p>No further construction is necessary or appropriate.</p>

a. Densify’s Opening Position

(i) VMware’s Proposed Construction Changes the Meaning of the Claim.

When construing a patent, “the name of the game is the claim.” *In re Hinker Co.*, 150 F.3d 1362, 1369 (Fed. Cir. 1998) (quoting Giles Sutherland Rich, *Extent of Protection and Interpretation of Claims—American Perspectives*, 21 Int’l Rev. Indus. Prop. & Copyright L. 497, 499 (1990)). As applied here, VMware’s proposed construction adds and deletes key words that would change the plain meaning. There is a reason VMware did not offer this construction for the 2019 *Markman* hearing (see D.I. 208), or during the second round of claim construction in December 2019 (see D.I. 426), or even at trial. It is inconsistent with the claim language and

cannot be supported by the intrinsic record. The Court already adopted VMware's earlier construction. *See* D.I. 356. VMware's newest construction is neither appropriate nor necessary.

(1) The Plain Language Forbids VMware's Construction.

VMware attempts to read the word “sets” out of the evaluating step, substituting “rules” for “rule sets.” It then goes further by modifying the language to state that every VM must have such “rules” pertaining to “each of” technical, business and workload constraints “associated with the corresponding virtual machine.” In other words, VMware's interpretation is that every VM must have rules associated with them and these rules must be all three of technical, business, and workload constraints. But the plain language makes clear that “rule sets” need not belong to VMs hosted within a virtualized environment — “data sets” do, as shown in claim 7 (emphasis added):

A method for validating an existing virtualized environment comprising a plurality of virtual machines placed on one or more virtual hosts, said method comprising: [a)] obtaining a data set for each of said plurality of virtual machines, **each data set** comprising *information pertaining to technical, business and workload constraints associated with a corresponding virtual machine*; [b)] evaluating the placement of said virtual machines in said virtualized environment using **said data sets** by evaluating each virtual guest against each virtual host and other virtual guests using one or more **rule sets** pertaining to said technical, business and workload constraints to determine guest-host placements; and [c)] identifying the existence of virtual machines with suboptimal placements to enable alternative placements for said virtual machines.

There are two “sets” discussed in claim 7: “data sets” and “rule sets.” “Data sets” contain “*information* pertaining to technical, business and workload constraints associated with a corresponding virtual machine.” For example, information pertaining to a business constraint might include the operating system. By contrast, “rule sets” are used to evaluate the “data sets” for each VM and host in the virtualized environment, and “one or more” of the rule sets must pertain to technical, business, and workload constraints. Using the same example, a rule set

might include a business rule saying Oracle VMs should be placed on Oracle hosts to minimize licensing fees. Those “rule sets” need not belong to the individual VMs and are, instead, used to evaluate each VM against all VMs and hosts in the virtualized environment. Rule sets typically reside at the virtualized environment (e.g., cluster) level and are applied to the VMs within that environment.

Indeed, it does not even make logical sense to require an individual VM to have a “rule set” because rule sets serve no purpose in isolation. Rule sets make sense in the context of a virtualized environment in which information pertaining to technical, business, and workload constraints must be evaluated, *i.e.* compared between at least two things, to determine the optimal placement of VMs on hosts. For example, the rule that an Oracle VM should be placed on an Oracle host only makes sense if there are at least two things to evaluate against each other. “Data sets,” on the other hand, contain information about individual VMs — *i.e.*, operating system, CPU, memory, maintenance window, etc. In our example, it makes sense that a single VM would contain data, *i.e.* its operating system, which pertains to a business rule. “Rule sets” are tools of comparison *between* VMs (and hosts), meant to permit evaluation of placement choices for VMs in a virtual environment. Claim 7 makes little sense if each VM in the virtualized environment must carry “one or more rule sets.” In other words, VMware’s construction is “not remotely close to what the claim says.” *Bio-Rad Labs., Inc. v. Int’l Trade Comm’n*, 998 F.3d 1320, 1331 (Fed. Cir. 2021). “Inventors are masters of their claims, and the words they use to describe and claim their invention are decisive and binding.” *Id.* The Federal Circuit “has repeatedly and clearly held that it will not read unstated limitations into claim language.” *Northern Telecom Ltd. v. Samsung Elecs. Co., Ltd.*, 215 F.3d 1281, 1290 (Fed. Cir.

2000). Here, the inventors made clear that “data sets” belong to each VM and “rule sets” are used to evaluate each VM against other VMs and hosts in the virtualized environment.

VMware also attempts to import limitations into the “evaluating” step that are inconsistent with the “obtaining” step. (D.I. 767 at 2). The prior “obtaining” step requires only that “information pertaining to technical, business and workload constraints associated with a corresponding virtual machine” is obtained. In other words, in the obtaining step, the invention collects information pertaining to the three constraints from a VM, and that then becomes a “data set” corresponding to that VM. This is not part of the “evaluating” limitation that VMware seeks to have construed. There is no basis to import limitations into the “evaluating” step that are already addressed in the “obtaining” step and are inconsistent with the “obtaining” step. In the evaluating step, the rule sets are used to evaluate the placement of VMs by evaluating the rule sets against each VM’s data set. And “one or more” of those rule sets must pertain to all three constraints.

Accordingly, the plain language of claims 3 and 7 does not support VMware’s construction.

(2) The Specification Doesn’t Require Rule Sets for Each VM.

The specification of the ’687 patent does not contain any language suggesting that every VM itself must have a rule set pertaining to technical, business, and workload constraints. Instead, the patent teaches that every VM has a data set with information pertaining to technical, business, and workload constraints (if any). The data set from each VM is then evaluated using one or more rule sets pertaining to technical, business, and workload constraints. As one example, the ’687 patent contemplates that a preferred embodiment allows obtaining “configuration compatibility results . . . for each rule set” and “for each source-target pair,” but

there is no such requirement for all embodiments that all VMs possess a tripartite rule set, only that each have *information* relevant to such a rule set — *i.e.*, operating system, CPU, memory, maintenance window, % CPU utilization, etc. ’687 pat. at 13:54-60. As a further example, the specification expressly contemplates that placement of some VMs will not be restricted by business rules, but that all VMs in the virtualized environment are evaluated as to all rule sets. *Id.* at 10:56-57; 13:18-45; 20:46-21:7. For example, a rule restricting placement of Oracle VMs to Oracle hosts would not restrict placement of a Linux VM. Nevertheless, that Linux VM would be “evaluated” to determine whether it satisfied the rule that Oracle VMs should go on Oracle hosts (it would not). Another embodiment envisions obtaining workload data in the first stage and then using that data in a second stage to determine technical- and business-related compatibility according to rule sets in a virtualized environment. *See* ’687 pat. at 11:5-22. Indeed, the embodiments do not mandate technical, business, or workload *rules* on any VM. As described above, that does not even make sense. Rules are comparison tools used to evaluate placement of VMs based on information in data sets. *See* ’687 pat. at 13:21-24 (“For each rule in the set **28**, the data referenced by the rule is obtained for both the target and source. The rule is evaluated by having the rule engine **90** compare the data.”). VMware’s proposed construction requiring every VM to have a business rule would “exclude[] all disclosed embodiments” in the specification of the ’687 patent, which “is especially disfavored” under Federal Circuit law. *Kaneka Corp. v. Xiamen Kingdomway Grp. Co.*, 790 F.3d 1298, 1304 (Fed. Cir. 2015). For all these reasons, the specification also does not support VMware’s proposed construction.

(ii) This Term Has Already Been Fully Construed and Is Law of the Case.

The Court has already construed this term. The Court should not revisit claim construction on the ’687 patent as VMware proposes. VMware’s attempt at a “do-over” violates

the law of the case doctrine.¹ *Waldorf v. Shuta*, 142 F.3d 601, 616 n.4 (3d Cir. 1998). “[A] fundamental precept of common-law adjudication is that an issue once determined by a competent court is conclusive.” *Arizona v. California*, 460 U.S. 605, 619 (1983). VMware had every opportunity to put forward its construction in September 2019 and did so. *See* D.I. 208, Ex. A at 1. Indeed, the Court *adopted* VMware’s construction. D.I. 357 at 1; *see also* D.I. 526 at 22. That should settle the matter.

VMware is, in substance, moving for reconsideration. But, as this Court has stated (*see* D.I. 946 at 3-4 n.2), the purpose of such a motion is to “correct manifest errors of law or fact or to present newly discovered evidence.” *Max’s Seafood Café ex rel. Lou-Ann, Inc. v. Quinteros*, 176 F.3d 669, 677 (3d Cir. 1999). VMware has not even attempted to meet this standard, nor could it. VMware cannot point to any error in the Court’s original construction because the Court’s construction is *VMware’s* construction. *See Nuance Commc’ns*, 813 F.3d at 1373.

Instead, VMware asserts that *O2 Micro International v. Beyond Innovation Technology Company*, 521 F.3d 1351 (Fed. Cir. 2008), supports the idea that district courts must resolve any and every “dispute” about a claim, regardless of when the dispute arose and without any traditional showing of need. D.I. 962 at 13; D.I. 968 at 13; D.I. 1003 at 13. In other words, VMware argues that *O2 Micro* created a special patent rule that exempts claim construction from the law that parties are bound by earlier decisions unless they can offer a compelling reason for reconsideration. But *O2 Micro* creates no such rule. It does not permit a party “dissatisfied with

¹ In its post-trial brief, VMware introduced for the first time, a new interpretation of the Court’s November 26, 2019 construction of this term — a construction that VMware had proposed (D.I. 208, Ex. A at 1), and that the Court adopted (D.I. 357 at 1). VMware’s proposed construction here attempts to rewrite the ordered construction now that it has come up with a new interpretation.

its own proposed construction” to seek a new one. *Nuance Commc’ns, Inc. v. ABBYY USA Software House, Inc.*, 813 F.3d 1368, 1373 (Fed. Cir. 2016).

Patents are not special in this regard. They are “legal instrument[s], to be construed, like other legal instruments, according to [their] tenor.” *Markman v. Westview Instruments, Inc.*, 517 U.S. 370, 388 (1996); *see also Teva Pharms. USA, Inc. v. Sandoz, Inc.*, 574 U.S. 318, 325 (2015) (“[A] judge, in construing a patent claim, is engaged in much of the same task as the judge would be in construing other written instruments, such as deeds, contracts, or tariffs.”). Parties can no more compel district courts to re-interpret the same provisions of a patent any more than they can the clauses of a contract or deed. The Scheduling Order set claim construction for the ’687 patent in 2019. D.I. 174 at 10. Two years and a jury trial have passed. “[G]ood cause” is required for a redo. *See* Fed. R. Civ. P. 16(b)(4); *see also* Fed. R. Civ. P. 60(b). None exists here.

VMware relies on the supposition that *O2 Micro* permits unlimited bites at the claim construction apple, whether or not they are justified. It does not. *O2 Micro* simply stands for the proposition that courts cannot avoid resolving timely disputes by forcing a jury to interpret patent claims. *See* 521 F.3d at 1360 (“When the parties raise an actual dispute regarding the proper scope of these claims, the court, not the jury, must resolve that dispute.”). In *O2 Micro*, the defendants offered competing constructions during the *Markman* hearing, but the Court refused to resolve that dispute and, instead, punted the question to the jury. *See id.* at 1359. That is not the case here, where the Court already adopted VMware’s proffered construction. *See Nuance Commc’ns*, 813 F.3d at 1373 (concluding there was no *O2 Micro* problem where the district court adopted the appellant’s proposed construction and the appellant “reversed course and tried to get a new construction of disputed terms shortly before trial, which the district court properly

denied given the parties’ earlier agreement and the lack of any good cause for revisiting the claim construction”); *see also, e.g., Function Media, L.L.C. v. Google, Inc.*, 708 F.3d 1310, 1327 (Fed. Cir. 2013); *Verizon Servs. Corp. v. Cox Fibernet Va., Inc.*, 602 F.3d 1325, 1334–35 (Fed. Cir. 2010); *Cordis Corp. v. Boston Sci. Corp.*, 561 F.3d 1319, 1338–39 (Fed. Cir. 2009). Indeed, interpreting *O2 Micro* as mandating reconstruction whenever a party comes up with a new idea to escape infringement would run afoul of the Supreme Court’s repeated admonition that patent cases are not governed by a set of “special rules.” *Trimble Inc. v. PerDiemCo LLC*, 997 F.3d 1147, 1154 (Fed. Cir. 2021); *see also Aatrix Software, Inc. v. Green Shades Software, Inc.*, 890 F.3d 1354, 1356 (Fed. Cir. 2018) (“Though we are a court of special jurisdiction, we are not free to create specialized rules for patent law that contradict well-established, general legal principles.”).

The same Rule 16 principles that apply to every other area of law apply to patent cases. *See, e.g., LifeNet Health v. LifeCell Corp.*, 837 F.3d 1316, 1322 (Fed. Cir. 2016); *see also, e.g., Akamai Techs., Inc. v. Limelight Nets., Inc.*, 805 F.3d 1368, 1376 (Fed. Cir. 2015); *Finjan, Inc. v. Secure Computing Corp.*, 626 F.3d 1197, 1207 (Fed. Cir. 2010). Rule 16 says parties must demonstrate “good cause” to deviate from the scheduling order. Rule 60(b) has similar requirements for relief from a court order. *O2 Micro* does not change that. *See Nuance Commc’ns*, 813 F.3d at 1373 (concluding the district court “properly denied” the appellant’s attempt “to get a new construction of disputed terms shortly before trial” given the “lack of any good cause for revisiting the claim construction”). There are no new circumstances that have created the need for a new claim construction. All the facts were available to VMware at the time of the first claim construction. The only thing that changed is the realization that Densify had proved infringement. VMware seeks to now change the claim language —by adding an

extra limitation — to provide a different defense. That it became dissatisfied with construction when it did not avoid infringement can hardly be a basis to revisit a claim construction.

Moreover, VMware’s argument that its new claim construction was a response to a purportedly new infringement theory — *i.e.*, that the addition of a single VM-host business rule infringes — does not reflect reality. *See* D.I. 767 at 3. It was VMware that was seeking to add a new limitation to the patent. Densify simply pointed this out. VMware offered a new (nonsensical) non-infringement theory in their post-trial supplemental brief, *i.e.*, that every VM must have technical, business, and workload rules. *See* D.I. 754 at 7 (“Any infringement theory that does not require business rules for each VM in a virtualized environment should be rejected.”). This new theory was inconsistent even with their opening post-trial brief, which required *evaluation* of each VM, not that every VM had business, technical, and workload rules. *See* D.I. 602 at 3 (“Restated, each VM in the existing virtualized environment must be evaluated using rule sets for all three constraints against (1) each host (the “Each-Host” evaluations) and (2) other VMs (the “Other-VM” evaluations).”). Densify simply pointed out that the plain language does not support VMware’s new construction and, given the way DRS operates, the addition of a VM-host business rule necessitates an infringement finding based on the plain language of the claims. *See* D.I. 755 at 3-7; D.I. 761 at 2-5.

The sequence of events shows that VMware was aware of all facts necessary to make its current claim construction argument at the time of the first claim construction, and therefore cannot constitute good cause for violating the scheduling order. Densify of course accused VMware’s product of infringing, and thus was pointing to a single VM-host business rule as infringing. VMware knew this because it knows precisely how its product works. It made no argument during claim construction that each VM had to have a business constraint or a business

rule associated with it. Indeed, even at trial it made no such argument. Only after trial did VMware argue that each VM must have a business rule or business constraint, as opposed to each VM being *evaluated* based on information pertaining to the constraints. Nothing changed about Densify’s infringement argument — it has always been that VMware’s product infringes. That VMware did not think to propose the construction it now advocates is not a problem for this Court to remedy under Rules 16(b)(4) and 60(b), not to mention well-established principles of finality and fairness. *See, e.g., Nuance Commc’ns*, 813 F.3d at 1373.

The rules against needless reconsideration of previous decisions are crafted so that the “course of ordinary litigation” can proceed. *Arizona*, 460 U.S. at 619. As applied here, a new construction may require reopening fact *and* expert discovery on infringement, validity, and damages — *i.e.*, every major issue in the case. To do so would be wasteful and manifestly unjust, especially when the principal issue justifying a new trial had nothing to do with claim construction — *i.e.*, the mere presence of VMware’s competitor. *See Kruse Tech. P’ship v. Volkswagen AG*, 544 Fed. App’x 943, 954 (Fed. Cir. 2013). The parties spent millions to prepare the ’687 claims for trial. That work should not be undone without good cause. *See Gindes v. United States*, 740 F.2d 947, 949 (Fed. Cir. 1984). VMware has offered none.

b. VMware’s Answering Position

Since the preliminary injunction, VMware has understood this term according to its plain meaning, which requires “evaluating each virtual guest . . . using one or more rule sets pertaining to said . . . constraints.” The claim language requires that every evaluation (*i.e.*, between each VM and each host, and between each VM and other VMs) uses, *in each case*, one or more rules for each of the three constraints associated with the VM being evaluated. As explained below, Cirba’s expert’s testimony and trial presentation show that it previously held the same understanding, which is also consistent with the specification.

Only in supplemental briefing *after* the post-trial hearing did Cirba first propose the “single-rule” construction that it now seeks.² According to Cirba’s new construction, as long as any constraint-related rule is used for just one evaluation of one VM against a host *or* another VM, the entire claim term (which requires multiple of such evaluations) is satisfied. (*See, e.g., supra* § II.C.1.a.(ii) at page 10 (alleging that “a single VM-host business rule [is] infringing,” and disagreeing that “each VM [must] have a business constraint” to infringe).) Cirba’s new position, which renders the three recited constraint types optional, is contrary to the plain language of the claim and the specification.³ It should be rejected.

The Court should not ignore the parties’ dispute as to the meaning of this limitation, as Cirba urges. Contrary to Cirba’s mischaracterizations, the Court has not previously construed this term, the scheduling order expressly calls for its construction, and VMware’s proposed construction is consistent with its previous non-infringement positions. The Court should address the differing interpretations that arose post-trial and resolve the “fundamental dispute regarding the scope of [this] claim term.” *O2 Micro Int’l, Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351, 1362 (Fed. Cir. 2008).

² Cirba does not seriously challenge that its claim construction has changed. Rather, it argues only that “[n]othing changed about Densify’s infringement argument [because its argument] has always been that VMware’s product infringes.” (*Supra* § II.C.1.a.(ii) at page 11.)

³ It is also seemingly inconsistent with Cirba’s amendment that added the disputed claim language to claims 3 and 7. (D.I. 1032-5 at 126-24 (amendment adding “by evaluating each virtual guest against . . . pertaining to said technical, business and workload constraints” to claims.) Cirba argued that its amendments distinguished prior art that only evaluated and “determine[d] placement of a single VM at a time.” (*Id.* at 133-34.) To the extent Cirba’s new claim construction position allows for infringement with evaluation and placement of a single VM, it is inconsistent with the prosecution history.

(i) The Intrinsic Evidence Supports VMware's Construction

The plain language of this limitation requires that, for each evaluation (whether between (1) a VM and a host or (2) a VM and another VM), one or more rules be assessed *in each case* for each of the three constraints. This is consistent with the requirement earlier in the claim that “technical, business, and workload constraints [are] associated with a corresponding virtual machine.” If the three constraints are associated with the VM being evaluated, the only way to “evaluat[e] each virtual guest . . . using one or more rule sets pertaining to said . . . constraints” is to assess constraints for each VM using constraint-related rules during each evaluation.

Although Cirba disputes whether all three constraints must be associated with each VM being evaluated, the intrinsic record demonstrates that there is no other reasonable interpretation. It confirms that constraint-related rules must be used in each VM evaluation against a host or another VM. If one evaluation between a VM and a host lacks a rule for any one of the three constraints (*e.g.*, there is no business constraints rule), the claim limitation is not met. The same is true for VM-against-VM evaluations.

(1) The Claim Language Supports VMware's Construction

The claim language requires each VM to have associated “technical, business, and workload constraints.” Element 7A expressly refers to “technical, business and workload constraints associated with a corresponding virtual machine.” As Cirba acknowledges (*supra* § II.C.1.a.(i).(1) at page 3), the phrase “*said* technical, business and workload constraints” in element 7B refers back to that phrase from element 7A (emphasis added):

[7A] obtaining a data set for each of said plurality of virtual machines, each data set comprising information pertaining to *technical, business and workload constraints associated with a corresponding virtual machine*;

[7B] evaluating the placement of said virtual machines in said virtualized environment using said data sets by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to *said technical, business and workload constraints* to determine guest-host placements;

Cirba argues, however, that “associated with a corresponding virtual machine” in element 7A modifies the “data set comprising information,” not the three constraints. (*Supra* § II.C.1.a.(i).(1) at page 3 (alleging that rules and constraints “need not belong to VMs . . . [but] ‘data sets’ do”); *id.* at 5 (alleging that information collected for a VM “becomes a ‘data set’ corresponding to that VM.”).) Cirba’s position ignores the syntactical structure of element 7A, which also recites “a data set for each of said plurality of virtual machines.” Cirba’s construction would render the phrase “associated with a corresponding virtual machine” in element 7A redundant.

Cirba’s position also ignores that “[m]odifiers should be placed next to the words they modify.” *HTC Corp. v. IPCom GmbH & Co., KG*, 667 F.3d 1270, 1274-75 (Fed. Cir. 2012). In element 7A, “associated with a corresponding virtual machine” appears next to “technical, business, and workload constraints.” It thus modifies the three constraints, not “data set.”

Contrary to its new position, Cirba’s trial presentation was consistent with VMware’s construction that the constraints (not the “data set”) are “associated with a corresponding” VM. After addressing the “data set” portion of element 7A, Cirba’s expert, Dr. Madisetti, separately addressed each of the three specific “constraints associated with a corresponding virtual machine.” He used highlighting to show that the constraints corresponded to VMs:

<p>Claim 7A</p> <p>obtaining a data set for each of said plurality of virtual machines, each data set comprising information pertaining to technical, business and workload constraints associated with a corresponding virtual machine;</p>	<p>Claim 7A</p> <p>obtaining a data set for each of said plurality of virtual machines, each data set comprising information pertaining to technical, business and workload constraints associated with a corresponding virtual machine;</p>	<p>Claim 7A</p> <p>obtaining a data set for each of said plurality of virtual machines, each data set comprising information pertaining to technical, business and workload constraints associated with a corresponding virtual machine;</p>	<p>Claim 7A</p> <p>obtaining a data set for each of said plurality of virtual machines, each data set comprising information pertaining to technical, business and workload constraints associated with a corresponding virtual machine;</p>
---	---	---	---

(D.I. 707 at 172-74, 176 (demonstrative highlighting showing constraints correspond to VMs).)

The claim language also requires that each VM’s evaluations (*i.e.*, against each virtual host and other VMs) use one or more rules for each of the three constraint types associated with the VM. That is the only way that each VM can be evaluated against each host and against other VMs “using one or more rule sets pertaining to *said* technical, business and workload constraints.” Each evaluation must use rules pertaining to the three constraints in each case, as (1) each VM is evaluated, (2) any evaluation uses rules pertaining to constraints, and (3) constraints are associated with the VM being evaluated.

The rest of the claim reinforces these strict requirements for the claimed evaluations. Element 7A requires, and Cirba concedes, that “a data set [must be obtained] for each of said plurality of virtual machines.” (*See supra* § II.C.1.a at page 5 (“‘data sets’ belong to each VM”).) To evaluate these data sets obtained for each VM, “one or more rule sets pertaining to said [associated] technical, business and workload constraints” are used. The claim term is not satisfied by only one VM evaluation using one or more rules for only one constraint type, as Cirba alleges. Instead, *each* VM evaluation must use one or more rules for each of the three constraints. The phrase “in each case” in VMware’s construction makes this clear.

(2) The Specification Supports VMware's Construction

The specification also supports the requirement that all VMs be evaluated and that each evaluation use one or more rules for each of the three constraints. The specification describes a “system consolidation analysis [that] computes the compatibility of a set of systems 16 *based not only on technical and workload constraints as exemplified above, but also business constraints.*” (’687 patent at 12:45-48 (emphasis added).) It “determines the optimal layout of the source systems onto the target systems based on technical, business *and* workload constraints according to a multi-dimensional compatibility and consolidation analysis.” (*Id.* at 6:61-65 (emphasis added); *see also id.* at 6:35, 19:32-35, 23:24-32, 25:3-9, 25:32-37, 27:40-46, 27:63-28:2, 30:29-32, 32:40-46, 34:61-35:29 (all confirming all three constraints are used in analysis).) As these passages confirm, the evaluations in the ’687 patent account for all constraint types for all VMs, not merely application of a single rule for one constraint for a single VM as Cirba now argues.

Moreover, when describing evaluating a VM against a host, the specification confirms that no VM-host evaluations are skipped, and no rules are skipped for any evaluation. Figure 10 (excerpted and annotated with highlighting below) illustrates that “each rule in [a] rule set” is evaluated “by comparing source and target data,” which correspond to a VM’s and host’s data set. (’687 patent at Fig. 10, 13:3-17, 13:46-67; *see also id.* at 13:18-45, Figs. 9, 11, 17.)

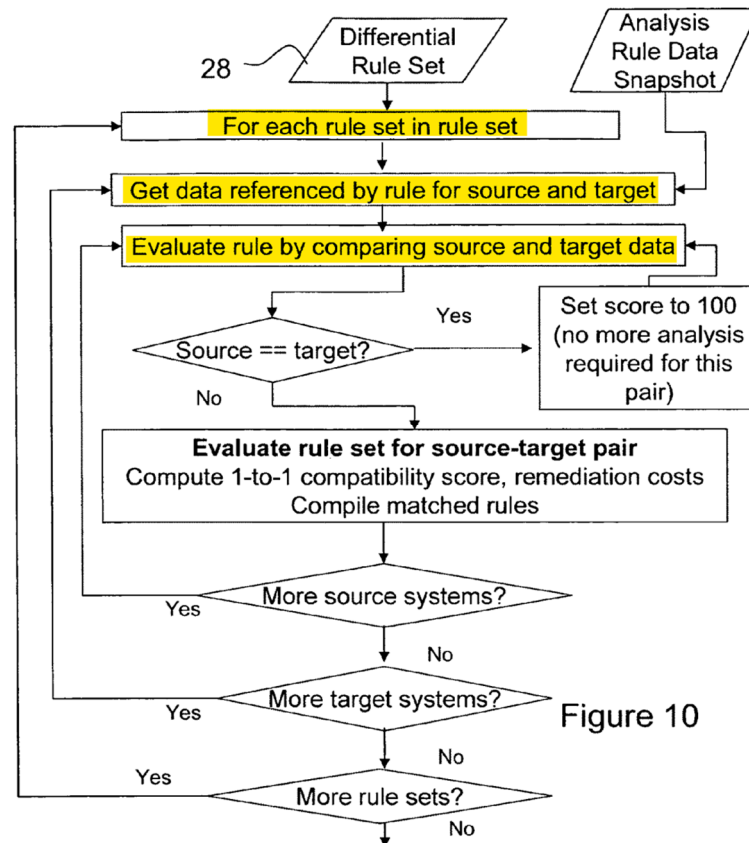


Figure 10

“The evaluation of the rules results in compatibility scores” for each VM-host pair that are “used for subsequent analyses, such as combining with workload compatibility results to obtain overall compatibility scores.” (*Id.* at 12:40-44.) “Compatibility maps” then display the resulting compatibility scores for *every* combination of systems. (*See, e.g., id.* at Figs. 30-32.) No evaluations between two systems are missing from the score maps (*e.g.*, with an “N/A” entry), and the specification never describes skipping rules in the evaluations.

And when describing evaluating a VM against another VM, the specification confirms that no VM-VM evaluations are skipped, and no rules are skipped for any evaluation. Annotated Figure 21 below shows that the same rule set is used for “each source” (*i.e.*, each VM) of a transfer set being evaluated with “each [other] source” (*i.e.*, other VM) of that transfer set. (’687 patent at 20:63-65.) All rules are used when evaluating each VM against another VM:

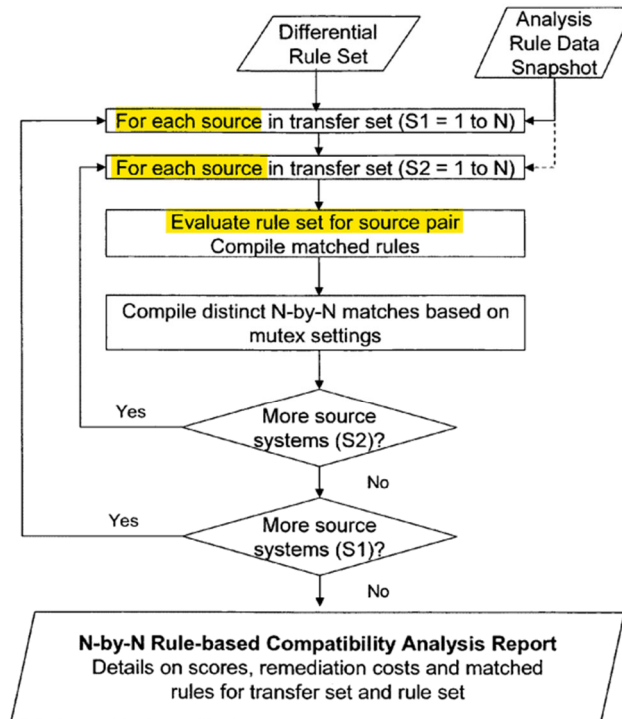


Figure 21

(ii) Cirba's Textual Counterarguments Lack Merit

Rather than providing its own express construction and explaining how the record supports it, Cirba focuses on attacking VMware's construction. But contrary to Cirba's argument, VMware does not read "rule sets" out of the claim. (*Supra* § II.C.1.a at page 3.) VMware's construction recognizes that term by referring to "one or more rules." This is consistent with the patent's description of "rule sets" as "groupings of rules." ('687 patent at 12:1-4.)

Further, Cirba's argument that VMs need not have all three technical, business, and workload constraints is inconsistent with the claim language. (*See supra* § II.C.1.a at page 5 ("[T]he patent teaches that every VM has a data set with information pertaining to technical, business, and workload constraints (*if any*).") (emphasis added).) As explained above, the plain

language requires that all three constraints be associated with each VM. None of Cirba's cited specification excerpts (which it fails to explain (*supra* § II.C.1.a at page 5-6)) is to the contrary:

- At 13:54-60, the specification explains that each VM is evaluated so that "compatibility results are [] generated for each rule set." This excerpt does not describe optional constraints or optional rules.
- At 10:56-57, the specification merely describes "evaluat[ing] the compatibility of every specified system as source-target pairs on a 1-to-1 basis." This excerpt does not mention rules or constraints at all.
- At 13:18-45, the specification confirms that each VM is evaluated using each rule set. "For each rule in the set 28, the data referenced by the rule is obtained for both the target and source. The rule is evaluated by having the rule engine 90 compare the data." ('687 patent at 13:21-24.) Nothing in this excerpt overrides the claim requirement that each VM has corresponding technical, business, and workload constraints and that rules are evaluated per each of these constraints.
- At 20:46-21:7, the specification confirms that "each source entity [is evaluated] against the other source entities with the rule set." (*Id.* at 20:64-65.) Nothing in this excerpt suggests that constraints or rules pertaining to those constraints are optional.

Next, relying upon VMware's non-infringement positions for Cirba's specific infringement allegations, Cirba contends that VMware's claim construction requires rules that "belong to" or be "on" VMs. (*See supra* § II.C.1.a at pages 3, 5; *accord supra* § II.C.1.a at pages 4, 6 (interpreting VMware's construction as requiring that a VM "possess" or "carry" rules).) Not so. Consistent with the plain claim language, VMware's construction requires only that that the recited "constraints" be associated with VMs and that "rule sets" (*i.e.*, one or more rules) pertain to each of a VM's associated constraints. For each evaluation (*i.e.*, of a VM against a host, and a VM against another VM), the one or more rules must be used to assess the VM's data set for each of the three constraints associated with the VM.

VMware’s non-infringement position is consistent with its proposed claim construction. Cirba’s infringement theory relied on DRS rules, which are assigned to specific VMs.⁴ To satisfy the requirement that each VM was evaluated against each host and other VMs using rules pertaining to all three constraints, “Cirba needed to show that *each* VM in the cluster had DRS rules so that it was evaluated (1) against each host in the cluster using technical, business, and workload constraints and (2) against other VMs in the cluster using the same three constraints.” (D.I. 754 at 5.) Cirba previously agreed, conceding it “needs to show [DRS rule] use with ‘every VM’ . . . in a [cluster].” (D.I. 672 at 7.) It does not follow from VMware’s non-infringement position that its *claim construction* requires rules associated with each VM. It only means that, *under Cirba’s specific infringement theory*, DRS rules needed to be specified for each VM.

Finally, Cirba’s Oracle VM to Oracle host infringement hypothetical both is inapt and does not support Cirba’s broad construction. (*See supra* § II.C.1.a at page 6.) According to Cirba, if a virtual environment has a single “business” rule requiring that Oracle VMs be placed on Oracle hosts, this would satisfy the “evaluating” step. Cirba deduces that this rule would be “evaluated” even against non-Oracle VMs (*e.g.*, a Linux VM) when determining if the rule applies. Cirba overlooks that the claims require “evaluating each virtual guest *against each virtual host and other virtual guests*.” In Cirba’s hypothetical, no evaluation against all “hosts” or “other virtual guests” would ever occur. The rule would be evaluated only for those specific VMs designated as Oracle VMs to determine whether they should be moved to specific hosts designated as Oracle hosts. No other virtual hosts or other VMs are involved.

⁴ DRS rules are created and assigned for specific VMs. (PTX-2057 at 98 (DRS rule created for “individual [VMs] that you select”).) A DRS rule for one VM does not apply automatically to others. Accordingly, to prove infringement by DRS rules, Cirba would need to show that each VM in a virtualized environment has at least one DRS rule, so that all required evaluations occur.

Regardless, Cirba’s infringement hypothetical also is neither disclosed by nor tethered to anything in the intrinsic or extrinsic evidence and thus cannot serve as a basis for broadening the claim scope. Cirba’s lone specification citation provides no support for its hypothetical. (*Supra* § II.C.1.a at page 6 (citing ’687 patent at 11:5-22).) The excerpt merely states that rules for technical, business, and workload constraints may be evaluated in different stages; it does not contradict the requirement that all three constraints be evaluated on a per-VM basis. In fact, it concludes by noting that the technical, business, and workload results are used for “overall compatibility scores [that] are computed for each pair of systems.” (’687 patent at 11:5-17.) This confirms that all three constraints are present in every VM evaluation against a host or another VM.

(iii) Cirba’s Procedural Arguments Also Lack Merit

Cirba’s contention that “[t]he Court has already construed this term” ignores what was actually construed previously and its own changing position. (*Supra* § II.C.1.a at page 6.) During the 2019 claim construction proceedings, the parties briefed and the Court construed the meaning of the phrase “evaluating each virtual guest against each virtual host and other virtual guests.” The parties’ and the Court’s focus was which “virtual hosts” needed to be evaluated with “each virtual guest.” (D.I. 356 at 5 (Court noting that dispute is about whether “‘each virtual guest’ must be evaluated against each virtual host in the virtualized environment (and not only against a subset of those virtual hosts).”).)

The disputed phrase now is “evaluating each virtual guest against each virtual host and other virtual guests *using one or more rule sets pertaining to said technical, business and workload constraints.*” The dispute arises from Cirba’s new interpretation of the second part of the term (emphasized), which was not at issue during the 2019 claim construction proceedings.

Cirba’s contention that VMware has deviated from the scheduling order by raising this claim construction dispute also is incorrect. (*Supra* § II.C.1.a at page 9.) Cirba ignores the Court’s express adoption of VMware’s proposal to resolve this claim construction dispute now. (D.I. 1003 ¶ 13 (“Further claim construction under this Scheduling Order shall [include] . . . the claim construction dispute for the ’687 patent that Cirba first raised in the parties’ supplemental post-trial briefing”); D.I. 968 at 13.) The Court rejected Cirba’s attempt to exclude the ’687 patent from the current claim construction process. (*See* D.I. 1003 ¶ 13 (adopting VMware’s proposal regarding claim construction for the ’687 patent and rejecting Cirba’s proposal to exclude it).

Cirba’s final argument—that VMware has changed its claim construction position—misstates the record and does not withstand scrutiny. (*See supra* § II.C.1.a at page 10 (alleging that it has always “point[ed] to a single VM-host business rule as infringing”).) To the contrary, it is Cirba that has changed its position. During the preliminary injunction proceedings, Cirba’s expert, Dr. Madisetti, made clear that Cirba’s infringement position required DRS rules “for each VM in a cluster of multiple hosts.” (D.I. 114, Ex. A at 205:10-206:9.) In its JMOL opposition, Cirba then contended that infringement required DRS rules for “every VM” in the virtualized environment. (D.I. 672 at 7.) And at the JMOL hearing itself, Cirba insisted that infringement requires “at a minimum” two types of DRS rules—“VM-to-VM and VM-to-host rules”—to satisfy the recited evaluations involving technical and business constraints. (D.I. 951 at 39:15-25.) Not until its request for supplemental post-trial briefing did Cirba argue that infringement only required that “a business rule [be] enabled” for a single VM. (D.I. 745 at 1.)

c. Densify’s Reply Position

VMware contends that its new construction reflects the plain meaning. *Supra* § II.C.1.b at page 11. But VMware proposed a different construction for the same claim term, the Court

adopted VMware's prior proposed construction, and it was VMware—and not Densify—that attempted to inject a new claim construction dispute into the post-trial proceedings after failing to prove noninfringement based on its prior claim construction. VMware's "do-over" construction is contrary to the intrinsic evidence, unsupported by the law, and should be rejected.

The overarching problem is that VMware fundamentally misapprehends how the claims work. In the obtaining step of claim 7, a "data set" is obtained from each VM. That "data set" is made up of (i.e., "comprises") information relating to (i.e., "pertaining to") technical, business and workload constraints associated with the VM. Then, in the evaluating step, "rule sets" are used to evaluate the placement of VMs by evaluating the "rule sets" against each VM's data set. And "one or more" of those "rule sets" must pertain to all three constraints. There is no requirement that "rule sets" are associated with particular VMs, and indeed this is not how the embodiments work. "Rule sets" typically are at the virtual environment (cluster) level. The "rule sets" are used to evaluate the "data sets" that are associated with each VM. And the claim language does not require that all "rule sets" have all three constraints – it plainly states that one or more rule sets must have all three constraints. This operation is supported by both the plain language of the claims, and is throughout the specification. VMware's interpretation seeks to change this plain language.

(i) VMware's "Do-Over" Construction is Contrary to the Intrinsic Evidence.

VMware's new construction first attempts to read the word "sets" out of the evaluating step—substituting "rules" for "rule sets"—and then adding an additional requirement that every VM must have such "rules" associated with them pertaining to "each of" technical, business, and workload constraints "associated with the corresponding virtual machine." As explained below, the plain language of the claims is contrary to VMware's new proposed construction. Nor does

the specification require that every VM itself (as opposed to the virtualized environment or cluster) must have a “rule set” pertaining to technical, business, and workload constraints—let alone “rules” pertaining to each. In other words, the claim speaks in terms of “rule sets,” not “rules.” And further, there is no requirement that a “rule set” is associated with a particular VM. “Data sets” are associated with VMs, not “rule sets.”

(ii) The Plain Language Contradicts VMware’s “Do-Over” Construction.

VMware’s argument boils down to trying to import language from 7A into 7B. Its view is because “data sets” are associated with particular VMs, “rule sets” must be also. *Supra* § II.C.1.b(i)(1) at pages 13-15. Element 7A (the “obtaining” step), first discloses “obtaining a data set for each of said plurality of virtual machines,” further disclosing that “each data set” comprises “information pertaining to [(i.e., relating to)] technical, business and workload constraints associated with a corresponding virtual machine.” Element 7B (the “evaluating” step), then discloses evaluating the placement of virtual machines using “one or more rule sets pertaining to said technical, business and workload constraints....” 7B uses the “rule sets” to evaluate the obtained “data sets.”

VMware’s argument seems to be that Densify’s understanding of the plain meaning of the claim language would make the term “associated with a corresponding virtual machine” in 7A somehow redundant with the term “a data set for each of said plurality of virtual machines” in 7A. *Supra* § II.C.1.b(i)(1) at page 14. It would not. The claim language simply says that 7A first requires obtaining data sets for the virtual machines. Then, the claim describes that each data set will also include information that pertains to the constraints, and to which virtual machine the information belongs to (is associated with). VMware’s citation of a handful slides from Dr. Madisetti’s trial demonstratives without context does them no favors. *See id.* As a

threshold matter, it is difficult to understand how highlighting of claim language on a slide makes its point. More substantively, it implies that the slides say that constraints are associated with the corresponding VM, not the data set. First, VMware misunderstands the use of “constraints.” Information pertaining to constraints associated with a VM are what comprises the “data set,” clearly in 7A. Then in 7B, “rule sets” pertain to the three constraints – i.e., there are rules sets related to technical, business, and workload constraints. Nothing cited by VMware in Dr. Madisetti’s presentation contradicts this plain language understanding. He was clear that “data sets” are information pertaining to technical, business, and workload constraints associated with a VM. *See* Ex. B-5 (Trial Tr.) at 475:3-12; 478:11-24; 480:4-19.

Claim 3, not addressed by VMware, illustrates the absurdity of its supposed “plain language” argument. Claim 3’s preamble and analyzing step makes clear that it is the rule sets that pertain to the constraints, not the VMs:

3. A method for performing a virtual to virtual (V2V) transformation for a plurality of existing virtual guests and hosts, said method comprising:

analyzing said existing virtual guests and hosts based on technical, business and workload constraints by evaluating each virtual guest against each virtual host and other virtual guests using one or more rule sets pertaining to said technical, business and workload constraints to determine guest-host placements;

Claim 3 (emphasis added). The same claim language in 7B has the same meaning. Claim 3 confirms the claimed evaluation uses one or more rules sets that pertain to technical, business, and workload constraints. VMware’s own expert confirmed at trial that claim 3 has the “same [e]valuating each to each” as claim 7, (Ex. B-6 Trial Tr. at 1342), and he used the “identical analysis” for the analyzing step of claim 3 and claim 7. *Id.* at 1357. VMware’s attempt to tie the constraints to individual VMs is nothing more than an attempt to inject a new limitation into claim 7, and would require the same term to mean two different things.

VMware's arguments in support of its new theory that "rule sets" (or apparently "rules") must exist on each VM rely on obfuscation, contending that Densify has taken a new position that "constraints" are associated with corresponding VMs. *Supra* § II.C.1.b(i)(1) at 14. As set forth above, Densify's position is neither new, nor what VMware attempts to characterize it as. The claim language is clear in 7A that "data sets" comprised of information pertaining to the three constraints are what is associated with the VMs. Densify's position has consistently been that "rule sets" typically reside at the virtualized environment (e.g., cluster) level and are applied to the VMs within that virtualized environment. *See* Ex. B-7 (2019 10 31 Hillier Dep. Tr.) at 476:9-477:12. VMware ignores this critical fact.

Similarly, VMware's new argument that the claim requires evaluation by rule sets for each VM pertaining to all three constraints in each case, improperly attempts to import limitations from the obtaining step (7A) into the evaluating step (7B). VMware's belated noninfringement-driven claim construction theory grossly ignores the "one or more" claim language immediately preceding "rule sets" in 7B (and claim 3). The explicit language of both claim 7 and claim 3 only requires that "one or more" of the claimed rule sets of 7B must pertain to all three constraints.

(1) The Specification Demonstrates The Error in VMware's "Do-Over" Construction.

VMware's cited portions of the specification do not support its new theory that all rule sets must have all three constraints. *Supra* § II.C.1.b(i)(2) at pages 16-17. The cited passages merely reiterate the claim language that, for every VM, a data set is obtained with information pertaining to technical, business, and workload constraints (if any), which is then used to evaluate placement using one or more rule sets pertaining to said technical, business and workload constraints. The specification passages cited by VMware merely recognize that

generally, all three constraints factor into the analysis, but the specification is consistent with the claims in disclosing a hierarchy where “data sets” are first obtained and then evaluated using “rule sets.”

Further, a complete review of the specification demonstrates the error of VMware’s attempt to change the explicit claim language and substitute “rules” for “rule sets” in its new proposed construction, because the specification clearly delineates between “rules” and “rule sets,” and only “rule sets” appear in the claims. *See* ’687 pat., Fig. 9:

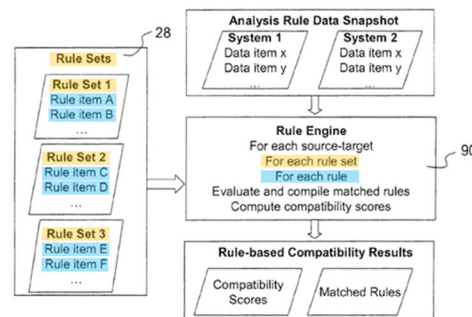


Figure 9

Indeed, the specification is replete with distinctions between “rule sets” and “rules.” *Id.* at 11:26-27 ; 12:6-8; 11:33-34; 12:1-4. “Rules” and “rule sets” are different.

In an apparent rehash of its non-infringement arguments from trial, VMware argues “the specification confirms that no VM-host evaluations are skipped, and no rules are skipped for any evaluation.” *Supra* § II.C.1.b(i)(2) at pages 17-18. But the specification makes clear that data sets—pertaining to technical, business, and workload constraints—are first obtained from each virtual machine in the virtualized environment, and then in the evaluating step, the rule set is used to evaluate the data set. *See id.* at 13:21-24. VMware cannot change the intrinsic evidence.

(2) VMware's Criticisms Of Densify's "Textual Counterarguments" Fail Because They Are Not Supported By The Plain Language.

VMware contends, without authority, that Densify is required to propose its own affirmative construction for a claim term that the Court has already construed when it adopted VMware's proposed construction. *Supra* § II.C.1.b(ii) at page 18. No construction is necessary when the claim language is clear. The Court's prior construction was the result of multiple briefs and oral argument. No further expenditure of the Court's resources is needed to change a construction that a jury can (and did) understand.

VMware also claims its new construction does not read "rule sets" out of the claim, but provides only a conclusory statement as to why it makes sense to replace "rule sets" with "one or more rules." *Id.* As described above, the specification clearly differentiates between "rules" and "rule sets." And VMware's argument "that all three constraints be associated with each VM" ignores key claim language. *Supra* § II.C.1.b(ii) at pages 18-19. The claims refer to "data sets"—not "constraints"—being associated with each VM. And the "data set" of each VM is comprised of "information pertaining to" technical, business, and workload constraints. The claims do not support VMware's arguments

Finally, VMware oversimplifies and distorts Densify's hypothetical concerning Oracle VMs. VMware seems to be saying Densify's position is that only a single "business" rule regarding placing Oracle VMs on Oracle hosts is required to satisfy 7B. *Supra* § II.C.1.b(ii) at page 20. This is not what Densify is saying. The point of the hypothetical was merely to demonstrate how a business rule in a rule set would be applied in the evaluating step, but Densify has always been clear that one or more rules pertaining to all three constraints are required to satisfy the claim. *See supra* § II.C.1.a(i)(2) at pages 5-6. It is difficult to understand how VMware then concludes that "[n]o other virtual hosts or other VMs are involved." Densify

accepts the Court’s prior claim construction that each VM is evaluated. But the question is *how* the claim “evalut[es] each virtual machine against each virtual host and other virtual machines.” (D.I. 357.) The claim language is clear the evaluation is accomplished using one or more rule sets pertaining to all three constraints.

(iii) VMware Fails To Offer Good Cause For Its Changed Construction.

VMware contends the Court *must* resolve its belated claim construction dispute created under *O2 Micro*. *Supra* § II.C.1.b at page 12. But the Federal Circuit has already squarely dealt with this issue and found that a party’s becoming “dissatisfied with its own proposed construction and [seeking] a new one does not give rise to an *O2 Micro* violation.” *Nuance Commc’ns, Inc. v. ABBYY USA Software House, Inc.*, 813 F.3d 1368, 1373 (Fed. Cir. 2016). The facts here are even more in favor of declining to consider VMware’s effective motion for reconsideration of its original proposed construction⁵. In *Nuance*, the appellant initially proposed the use of plain and ordinary meaning for the claim term at issue and the district court agreed, adopting its construction. Here, VMware did not attempt to simply rely on a plain and ordinary meaning construction, but rather affirmatively asked the Court to adopt its construction, which the Court did. The Court need not entertain here VMware’s attempt to “reverse[] course and tr[y] to get a new construction for [a] disputed term[],” where VMware has not demonstrated “any good cause for revisiting the claim construction.” *Id.* at 1373.

VMware attempts to shift blame, contending the “parties’ and the Court’s focus was which ‘virtual hosts’ needed to be evaluated with ‘each virtual guest.’” *Supra* § II.C.1.b(iii) at

⁵ VMware cannot point to a “manifest error of law or fact,” nor can it point to “newly discovered evidence” to justify a motion for reconsideration of a claim construction it prevailed on in the first instance. *See Max’s Seafood Café ex rel. Lou-Ann, Inc. v. Quinteros*, 176 F.3d 669, 677 (3d Cir. 1999).

page 21. But VMware chose its claim construction and noninfringement arguments. That it now regrets that choice does not provide an opportunity to re-do claim construction. VMware had months to develop its non-infringement theories, including through two 7-hour depositions of Andrew Hillier and a deposition of Tom Yuyitung, both named inventors of the '687 patent, and two 7-hour depositions of Densify's expert, Dr. Vijay Madiseti, not to mention extensive trial testimony and the assistance of its own expert. VMware did not ask questions directly about the "data sets" central to its new proposed construction in its deposition of Mr. Hillier. *See* Ex. B-7 at 467:2 (the sole mention of "data sets" in the transcript). VMware did not substantively explore the "rule sets" theory central to its new proposed construction in its depositions of Dr. Madiseti either. *See* Ex. B-8 at 51-53, 78-80, 132, 216, 219, 224-25 ("rule sets" mentioned 12 times total in December 2019 transcript). Indeed, VMware cites no evidence that its expert previously disclosed this "rule sets" theory in his report or at trial. VMware suggests no one was focused on the issue, *see Supra* § II.C.1.b(iii) at page 21, but that is incorrect. This was not an issue because it simply is not a limitation in the claim.

VMware argues this dispute is a result of Densify raising a claim construction issue in supplemental post-trial briefing. *See Supra* § II.C.1.b(iii) at page 22. Densify did no such thing. Densify proved its case based on the plain claim language, so now VMware wants to add a new limitation – that every VM must have a rule with all three constraints associated with it. After trial, VMware made a new non-infringement argument that each VM must have a business rule. Densify merely pointed out the claims require no such thing. That VMware sought to add a new limitation to the claims is not Densify raising a claim construction issue. It is akin to VMware now stating the claim requires the system run on a Macintosh computer. If Densify were to point out the claim requires no such thing, this can hardly be Densify raising an issue. This is not a

dispute about ambiguous words; this is VMware attempting to add a limitation simply not in the claims.

d. VMware's Sur-Reply Position

Cirba's Reply attacks a strawman construction that VMware never advanced. Cirba's lack of challenges to VMware's actual construction, the intrinsic support for that construction, and Cirba's inability to articulate its own construction show that VMware's position is correct.

(i) Each Evaluation Must Use Rules for Each Constraint

VMware's construction requires that "rule sets" pertaining to each of three constraints are used "in each case" of a VM evaluation against a host or a VM. It also requires that constraints, not rules, be associated with each VM being evaluated. Absent these clarifications, the jury will not know whether infringement requires (i) "constraints" for each VM and (ii) applying "rule sets" for all three constraints during each separate evaluation, or just collectively during all evaluations in the aggregate. Cirba's non-construction leaves these issues uncertain.

VMware's construction flows from the plain claim language. First, each of "technical, business, and workload constraints" must be associated with each "corresponding" VM. This requirement is clear from element 7A and the specification's description of the benefits of using all three constraint types. (*Supra* § II.C.1.b at pages 13-17.) If a VM has two associated constraints (*e.g.*, technical and workload) but not the third (*e.g.*, business), the claim is not met. Second, each evaluation of a VM against a host or another VM must use one or more rules for each of its three constraints. The specification confirms that rules are applied for every evaluation and are never missing. (*Supra* § II.C.1.b at pages 16-18.) If a single evaluation of a VM against a host or another VM skips a rule pertaining to even one of its three constraints, the claim is not met.

(ii) VMware's Construction Does Not Require Rules on Each VM

VMware's construction does not require that a rule set be associated with each VM, as Cirba claims. VMware's prior briefing (*supra* § II.C.1.b at page 19) could not have been clearer:

Cirba contends that VMware's claim construction requires rules that "belong to" or be "on" VMs. (*See supra* § II.C.1.a at pages 3, 5; *accord supra* § II.C.1.a at pages 4, 6 (interpreting VMware's construction as requiring that a VM "possess" or "carry" rules).) Not so. Consistent with the plain claim language, *VMware's construction requires only that that the recited "constraints" be associated with VMs* and that "rule sets" (i.e., one or more rules) pertain to each of a VM's associated constraints.

(Emphasis added.) Cirba's attacks on whether *rules* must be associated with VMs are irrelevant.

(iii) VMware's Construction Does Not Require that All Rule Sets Have Three Constraints

Nor does VMware require that "all rule sets must have all three constraints." (*Supra* § II.C.1.c at pages 14-15.) VMware's position is that, for each evaluation of a VM against a host or another VM, one or more rules must be assessed that collectively pertain to all of its three constraints. No constraints can be missing when all the rules (in aggregate) are assessed for an evaluation. VMware's construction does not require that a single rule set simultaneously pertain to all three constraints.

(iv) All Three Constraints Are Associated with VMs and Are Not Optional

Cirba cannot dispute that Claim 7 requires that "technical, business, and workload constraints [be] *associated* with a corresponding virtual machine" (emphasis added). Nor does Cirba dispute that its trial presentation was consistent with VMware's plain meaning interpretation. (*See* § II.C.1.b at page 15 (highlighting showing "constraints associated with a VM").)

Instead, Cirba argues that the claims require only that “‘data sets’—not ‘constraints’—be[] associated with each VM.” (*Supra* § II.C.1.c at page 28.) But Cirba cites no support for its position.

Cirba repeats its “if any” position that a VM need not have all three constraints available for evaluation. (*Supra* § II.C.1.c at page 26.) But it cites no support for this position, and the claim language makes clear that the three constraint types are not optional. Each VM’s data set must have information for three constraints associated with that VM.

(v) Cirba’s Distinction between “One or More Rules” and “One or More Rule Sets” Is a Distinction without a Difference

Per Cirba, a “rule set” cannot be “one or more rules,” as the specification refers to “rules” and “rule sets” separately. (*Supra* § II.C.1.c at pages 27, 28.) Cirba ignores that a “rule set” is simply “one or more rules.” Figure 9, which depicts three rule sets each having two rules, is consistent with VMware’s construction. Cirba does not explain how defining “one or more rule sets” to be “one or more rules” contradicts Figure 9 or would read “rule sets” out of the claims.

(vi) VMware Did Not Delay Raising the Need for Claim Construction

To avoid adoption of VMware’s construction, Cirba again argues it is untimely. Cirba again ignores that the scheduling order provides for construction of this term. (D.I. 1003 ¶ 13.)

Cirba is also wrong that VMware did not previously “substantively explore” non-infringement under its proposed claim construction. (*Supra* § II.C.1.c at page 30.) Cirba does not acknowledge, much less address, that VMware probed its non-use of “rules” when deposing Cirba’s technical expert. (*Supra* § II.C.1.c at page 22.) Cirba also ignores that VMware’s interrogatory responses explained VMware’s non-use of “rules” based on its plain meaning. (Ex. A-15 at 11-12). Cirba further ignores that VMware’s technical expert provided a report

explaining why VMware does not satisfy this limitation. (Ex. A-16 ¶¶ 89-96, 130-159.) Finally, Cirba overlooks that VMware cross-examined Cirba’s expert about “rules” at trial (TT at 618:1-20; D.I. 511 at DDX-3.3) and raised the issue throughout its post-trial briefing (D.I. 602 at 2-3; D.I. 685 at 1-4; D.I. 754 at 6). Tellingly, Cirba fails to cite anywhere in the record it previously pursued its current theories.

(vii) The Court Did Not Construe the Term Before and Should Do So Now

Cirba’s arguments that the Court has already construed this limitation and that VMware seeks a “new construction” or “do-over” is divorced from reality. As VMware explained (Resp. at 10), the 2019 claim construction efforts addressed a different limitation—one that did not include “rule sets” or “constraints. The Joint Claim Construction Statement (D.I. 208 at 1), parties’ briefing (D.I. 270 at 7-22), and Court order confirms this. (D.I. 356 at 4-6; D.I. 357.) “Rules sets” and “constraints,” by contrast, are the focus of the parties’ current dispute.

D. Terms Common to U.S. Patent No. 10,951,459 (“’459 Patent”) and U.S. Patent No. 10,523,492 (“’492 Patent”)

**1. “source system” / “source computer system”
(’492 patent: claims 1, 12, 23) (’459 patent: claims 1-3, 9-10, 16-17, 22-23, 26-34, 40-41, 47-48, 53-54, 57-63)**

VMware’s Proposed Construction	Densify’s Proposed Construction
“a system from which applications and/or data are to be moved”	“a physical, virtual, or hypothetical system from which applications and/or data are moved or are to be moved”

a. Densify’s Opening Position

The patentees acted as their own lexicographer for the term “source system.” The ’492 and ’459 patents state that “a *source system* refers to a system from which applications and/or data are to be moved, and a target server or system is a system to which such applications and/or

data are to be moved.” ’492 pat. at 5:56-59 (emphasis added); ’459 pat. at 5:57-60 (emphasis added). Both patents also broadly define “system,” explaining that “systems” may be “physical systems, virtual systems or hypothetical models.” ’492 pat. at 5:64-67; ’459 pat. at 5:65-6:1.

The patents also make clear that “source systems” include systems from which applications and/or data “are moved” or “are to be moved.” The patents explain that “a source system refers to a system from which applications and/or data are to be moved.” ’492 pat. at 5:56-63; ’459 pat. at 5:57-64. The patents also explain that “[c]onsolidation . . . can be considered to include one or more ‘transfers,’” where “[t]he *actual transfer* describes the *movement* of a single source entity onto a target” and “[t]he transfer type (or consolidation strategy) describes how a source entity *is transferred* onto a target, e.g. virtualization, OS stacking etc.” ’492 pat. at 6:34-40 (emphasis added); ’459 pat. at 6:35-41 (emphasis added). Thus, the patents describe “source systems” not only as systems from which applications and/or data “are to be moved” but also from which they “are moved,” e.g., “by moving” or through an “actual transfer.”

b. VMware’s Answering Position

VMware’s proposed construction follows the specification’s express definition: “a source system” is “a system from which applications and/or data are to be moved.” (’492 patent at 5:54-55.)⁶ VMware’s proposed construction also follows Cirba’s own prior positions.

Cirba has stated repeatedly that it acted as its own lexicographer. In its preliminary response to VMware’s IPR petition against the ’492 patent, Cirba argued that “[t]he ’492 patent’s specification *expressly defines* ‘source system’ to mean ‘a system from which applications and/or data are to be moved,’” the same express definition that VMware proposes as

⁶ The ’459 patent is a continuation of the ’492 patent and shares a common specification.

its construction now. (D.I. 1032-16, Ex. CC at 17 (emphasis added).) The PTAB adopted this construction when instituting the IPR. (*Id.*, Ex. DD at 8 (agreeing with Cirba’s express-definition argument, the PTAB “preliminarily construe[d] ‘source system’ to mean ‘a system from which applications and/or data are to be moved’”).)

And at the Section 101 hearing on the ’492 patent, Cirba confirmed this “definition”:

Ms. Hoang: The patent makes resoundingly clear, there is literally no exception[], this is a patent about computer system[s]. It’s in the title, it’s in the technical field, it’s in the overview. *And it is also in the lexicography of the patentee’s language. ‘Source system’ refers to systems in which applications and/or data are removed. The same is true of ‘target systems.’* You are talking about computers. . . .

Chief Judge Stark: . . . I take it, if and when we get there, you are going to argue that ‘system’ must be read as limited to [a] computer system. That is your implicit claim construction position; right?

Ms. Hoang: Yes, Your Honor. I think that is explicit in the lexicography.

(Ex. A-13 at 27-28 (emphasis added); *see also id.* at 12-13.)

Cirba now seeks both broadening and narrowing changes to the definition. First, without explaining what its addition actually means, Cirba proposes expanding the term’s scope by adding the phrase “are moved” as an alternative to the express definition’s use of “are to be moved.” Second, Cirba proposes that a source system be limited specifically to “physical, virtual, or hypothetical systems.” As VMware advances the identical construction that Cirba previously proposed in the ’492 patent IPR proceeding, its new construction should be rejected.

(i) The Inclusion of “Are Moved” Is Improper

Cirba has not explained why it seeks to add “are moved,” when the express definition includes only “are to be moved.” If “are moved” refers only to future moves of applications and/or data, it adds nothing beyond the “are to be moved” already included in the express

definition. On the other hand, if Cirba intends “are moved” to refer to past or present (real-time) moves, the specification provides no support. The ’492 and ’459 patents disclose only forward-looking “analysis” activity. The specification does not describe compatibility or consolidation analyses involving source systems that are actively being moved or have already been moved to a target system.

For example, the 1-to-1 compatibility analysis evaluates systems as both *potential* sources and *potential* targets, which precludes past moves and moves in real time. (’492 patent at 26:24-34.) The N-to-1 and N-by-N multi-dimensional compatibility analyses utilize “transfer sets” in which the one or more source systems are systems “*to be transferred* to the target” (*i.e.*, in the future). (*Id.* at 6:41-50, 29:59-67, 37:21-24 (emphasis added).) And the consolidation analysis is purely theoretical. During that analysis, transfer sets are added to an “intermediate consolidation solution” and removed from the pool of “available sources and targets.” The transfers are not actually executed, and no subsequent iteration evaluates source systems that have already been assigned. (*Id.* at 34:27-67, Fig. 30.)

Limiting the construction to the phrase “are to be moved,” as VMware proposes, is consistent with the specification’s express definition and correctly indicates that the purported invention identifies, but does not execute, system consolidations.

(ii) Cirba’s Proposed Limitation to “Physical, Virtual, or Hypothetical Systems” Should Be Rejected

The specification states that “systems” can broadly encompass “*any entity* which is capable of being analyzed for any type of compatibility.” (’492 patent at 6:28-32.) The specification expressly states, moreover, that a “system” “*should not be considered limited* to existing or hypothetical physical or virtual systems.” (*Id.* at 6:28-32 (emphasis added); *see also*

id. at 6:16-18 (“The compatibility analysis can also be generalized to evaluate entities beyond physical, virtual or hypothetical systems.”).)

Yet that is exactly what Cirba proposes. It would be legal error to limit the term in this manner. *See Vanderlande Indus. Nederland BV v. Int’l Trade Comm’n*, 366 F.3d 1311, 1320, 1322 (Fed. Cir. 2004) (declining to adopt narrow construction, where specification “expressly contemplate[d] other embodiments”).

Cirba’s addition also is likely to confuse the jury. It complicates rather than simplifies the term, as it adds words not found in the claim and does not replace the word “system” with clarifying language.

c. **Densify’s Reply Position**

Densify’s proposed construction is consistent with the specifications of the ’492 and ’459 patents, where patentees acted as their own lexicographer for the term “source system.”

VMware’s proposed construction leans heavily on a separate IPR proceeding, but provides no authority for why a construction adopted in a preliminary response before the PTAB should control. *See Pitney Bowes, Inc. v. Hewlett-Packard Co.*, 182 F.3d 1298, 1308 (Fed. Cir. 1999). Regardless, Densify’s proposed construction here is based on the same information it provided to the PTAB. Ex. B-9 (*VMware, Inc., v. Cirba IP Inc.*, IPR2021-00008, Paper 8 at 17-18 (P.T.A.B. Feb. 18, 2021).) VMware wants to capitalize on the specification express definition by willfully blinding itself to other disclosures in the specification. *Supra* § II.D.1.b at pages 35-36. Densify’s proposed construction reflects these other disclosures that are consistent with the express definition while providing more specificity. *See Trading Techs. Int’l, Inc. v. eSpeed, Inc.*, 595 F.3d 1340, 1353-56 (Fed. Cir. 2010) (affirming construction that “made two important changes to [an] express definition” in the specification based on additional context in “the claims, the rest of the specification, and the prosecution history”).

(i) The Specification Supports “Are Moved.”

The patents explicitly teach “[t]he [compatibility] analysis can be performed on an analysis with *pre-existing* source-target transfers.” ’492 pat., 33:32-34 (emphasis added); ’459 pat., 34:44-46 (emphasis added). The patents also teach “the compatibility analysis can evaluate the incremental effect of adding other source systems . . . to the specified transfer sets,” whereby a source system “can be individually assessed against [a] transfer set” that already includes one or more source systems stacked onto a target system. ’492 pat., 29:43-50; ’459 pat., 30:18-26. This demonstrates that the specification is not limited to “forward-looking ‘analysis’ activity,” but also compatibility analyses involving source systems that are *already* moved to a target system. Indeed, the patents teach the compatibility analysis may be used for purposes other than consolidation, including activities that can be performed *after* source systems have been moved to target systems, including “optimization,” “administration,” and “change.” ’492 pat., 38:13-18; ’492 pat., 9:43-48. VMware’s proposed construction ignores these teachings and should be rejected.

(ii) The Patents Disclose Physical, Virtual, and Hypothetical Systems.

The patents explain that “systems” may be “physical systems, virtual systems or hypothetical models.” ’492 pat., 5:64-65; ’492 pat., 5:65-66. VMware argues “systems” can be other things (*Supra* § II.D.1.b(ii) at pages 37-38 (citing ’492 pat., 6:16-18 and 6:28-32)) but its citations in support selectively omit key language disproving this construction. The specification teaches that “systems” include *components* of systems, but that they are still physical, virtual or hypothetical:

The compatibility analysis can also be generalized to evaluate entities beyond physical, virtual or hypothetical systems. **For example, entities can be components that comprise systems such as applications and database instances.** By analysing the

compatibility of database instances and database servers with database stacking rule sets, database consolidation can also be assessed. Similarly, application consolidation can be evaluated by analyzing application servers and instances with application stacking rules. The entity could also be a logical application system and technical data can pertain to functional aspects and specifications of the entity.

It will therefore be appreciated that a “system” or “computer system” hereinafter referred, can encompass any entity which is capable of being analysed for any type of compatibility and should not be considered limited to existing or hypothetical physical or virtual systems etc.

’492 pat., 6:16-32 (emphasis added); ’459 pat., 6:17-33 (emphasis added).

VMware’s quote from Densify’s argument during the § 101 hearing is off mark. Densify stated the patent’s lexicography requires that “‘system’ must be read as limited to [a] computer system.” *Supra* § II.D.1.b at page 36; *see also* 20-272-LPS, D.I. 97 at 3-4 (“adopt[ing] Densify’s implicit proposed claim construction, which is that ‘systems’ in claim 1 is limited to ‘computer system[s].’”). That is true and consistent with Densify’s proposed construction.

d. VMware’s Sur-Reply Position

(i) Including “Are Moved” Would Be Improper

In its Reply (*supra* § II.D.1.c at page 39), Cirba clarifies that its proposed addition of “are moved” refers to source systems that have been “*already* moved to a target system.” But the phrase “are moved” is simply confusing. On its face, the phrase does not mean “have already been moved” and could refer to present (*i.e.*, real-time) moves. (*Supra* § II.D.1.b at page 36-37.)

Regardless, neither “are moved” nor “already moved” is consistent with the patents’ purported invention or their express definition of “source system,” which Cirba previously adopted. (*Supra* § II.D.1.c at page 35-36.) The specification’s reference to “*pre-existing* source-target transfers” (’492 patent at 33:28-36) does not support Cirba’s new addition (*Supra* § II.D.1.c at page 39):

As part of a compatibility analysis input specification, systems can be designated for consideration as a source only, as a target only or as either a source or a target. These designations serve as constraints when defining transfers in the context of a compatibility analysis. *The analysis can be performed on an analysis with pre-existing source-target transfers.*

(Emphasis added.) As the full excerpt makes clear, “pre-existing source-target transfers” does not refer to sources already moved to targets, but to pre-assigning source to targets at the start of the analysis. (’492 patent at Fig. 29 (“source/target designation” as “compatibility analysis input”).)

Regardless, the sources in these pre-assigned transfers are not considered in the compatibility analysis, *i.e.*, they are not included in the claims’ “source systems.” (See *e.g.*, *id.* at 34:48-52 (“[E]ntities referenced by the transfer set are removed from the list of available sources and targets . . .”).) If sources had already been moved, as Cirba alleges, or pre-designated to be moved, the claims would not need to “evaluat[e]” and then “plac[e]” them again.

Cirba’s other excerpts underscore that compatibility analyses are forward-looking and do not relate to already-executed analyses. (*Id.* at 29:38-50 (“compatibility analysis can evaluate the incremental effect of adding other source systems”); *id.* at 35:12-52 (analyses performed on source and target “candidates”); *id.* at 3:3-6, 6:41-50, 7:35-38 (compatibility analysis involves “postulated transfer sets”; also describing “transfer sets” as proposed transfers or including source systems “capable of being transferred” to a proposed target); *id.* at Figs. 1-2, 9, 11, 14-15, 17, 21, 24(a)-(d), 25, 29 (flowcharts ending with “analysis,” “results,” or “determination”).) The claims, which focus on consolidation analysis, do not recite post-analysis activities such as “‘optimization,’ ‘administration,’ and ‘change.’” (*Supra* § II.D.1.c at page 39; see Ex. A-17 at 19:14-21, 21:20-22 (at § 101 hearing, Cirba agreeing that invention concerns analyzing system consolidation).)

**(ii) Limitation to “Physical, Virtual, or Hypothetical”
Should Be Rejected**

At the Section 101 hearing, Cirba represented that a “system” is a “computer system.” (*Id.*; Ex. A-17 at 12-13, 27-28.) Despite arguing that the claims are non-abstract on that basis, Cirba now argues “system” includes “hypothetical systems.” Worse, Cirba newly expands its construction to include “parts” of systems, such as “applications and/or data.” (*Supra* § II.D.1.c at page 39-40.)

The Court should reject Cirba’s contradictory and evolving additions and rely on the plain and ordinary meaning for “system.” Such a construction would be consistent with the testimony of Dr. Madisetti, who steadfastly refused to define “system.” (Ex. A-18 at 14:16-21:18, 51:11-53:14, 62:19-68:21 (during ’492 IPR deposition, repeatedly stating that “the specification speaks for itself” and that one should “defer to the specification” for its meaning).)

2. **“target system” / “target computer system”**
 (’492 patent: claims 1, 4-6, 9, 11-12, 15-17, 20, 22-23, 26-28, 31, 33)
 (’459 patent: claims 1-3, 9-10, 16-17, 22-23, 26-34, 40-41, 47-48, 53-54, 57-63)

VMware’s Proposed Construction	Densify’s Proposed Construction
“a system to which applications and/or data from a source system are to be moved”	“a physical, virtual, or hypothetical system to which applications and/or data are moved or are to be moved”

a. Densify’s Opening Position

The patentees also acted as their own lexicographer for the term “target system.” The ’492 and ’459 patents state that “a target server or system is a system to which such applications and/or data are to be moved.” ’492 pat. at 5:56-59; ’459 pat. at 5:57-60. And as with the source “system,” target “systems” may be “physical systems, virtual systems or hypothetical models” also applies to target systems. *See* ’492 pat. at 5:64-67; ’459 pat. at 5:65-6:1.

“Target systems” likewise include systems from which applications and/or data “are moved” or “are to be moved.” For instance, the patents explain that “a target server or system is a system to which such applications and/or data *are to be moved*.” ’492 pat. at 5:56-59 (emphasis added); ’459 pat. at 5:57-60 (emphasis added). As an example, the patents describe “an underutilized environment having two systems 16 [that] can be consolidated to a target system (one of the systems) *by moving* applications and/or data from the source system (the other of the systems) *to the target system*.” *Id.* (emphasis added). The patents also explain that “[c]onsolidation . . . can be considered to include one or more ‘transfers,’” where “[t]he *actual transfer* describes *the movement* of a single source entity *onto a target*” and “[t]he transfer type (or consolidation strategy) describes how a source entity *is transferred onto a target*, e.g. virtualization, OS stacking etc.” ’492 pat. at 6:34-40 (emphasis added); ’459 pat. at 6:35-41 (emphasis added). Thus, the patents describe “target systems” not only as systems to which applications and/or data “are to be moved” but also to which they “are moved,” e.g., “by moving” or through an “actual transfer.”

b. VMware’s Answering Position

Similar to “source system,” the ’492 and ’459 patents expressly define “target system” (and by extension, “target computer system”): “a target server or system is a system to which such applications and/or data [from a source system] are to be moved.” (’492 patent at 5:56-59.)

Here, too, VMware advances the identical construction that Cirba proposed in the ’492 patent IPR proceedings. Like the construction for “source system” above, Cirba proposed VMware’s construction in its preliminary response to the IPR, and the PTAB adopted it. (*See* D.I. 1032-16, Ex. CC at 18; *id.*, Ex. DD at 8.) For the same reasons explained above for “source system,” the correct construction is “a system to which applications and/or data from a source

system are to be moved.” Cirba’s insertions of “physical, virtual, or hypothetical” and “are moved” lack merit.

c. Densify’s Reply Position

The same arguments for “source (computer) system” apply to “target (computer) system.”

d. VMware’s Sur-Reply Position

For the same reasons as for “source system,” the patents expressly define “target system.” (’492 patent at 5:56-59.) Cirba’s additions to the express definition lack merit.

3. “place” terms (’492 patent: claims 1, 12, 23) (’459 patent: claims 1, 32, 63)⁷

Claim Terms (’492 patent: claims 1, 12, 23)	VMware’s Proposed Constructions
“for placing the source systems on target systems”	“for moving applications and/or data from source systems onto target systems”
“to determine whether the systems can or can not be placed together on a specific target system”	“to determine whether applications and/or data can or can not be moved from the systems together onto a specific target system”
“placing the source systems onto the target systems”	“moving applications and/or data from the source systems onto the target systems”
Claim Terms (’459 patent: claims 1, 32, 63)	VMware’s Proposed Constructions
“determining a placement of source computer systems on target computer systems”	“determining movement of applications and/or data from source computer systems onto target computer systems”
“determine a placement of at least one source system from the collection of computer systems on at least one target system from the collection of computer systems”	“determine a movement of applications and/or data from at least one source system from the collection of computer systems onto at least one target system from the collection of computer systems”
“one or more other source systems[,] either already placed on the specific target system, or	“one or more source systems from which applications and/or data[,] have either already moved onto the specific target system or are

⁷ Densify’s proposed construction for all “place” terms is “plain and ordinary meaning.”

being evaluated for placement onto the specific target system”	being evaluated for movement onto the specific target system”
“to determine if the specific source system can be placed with those other source systems on the specific target system”	“to determine if applications and/or data can be moved from the specific source system, with applications and/or data from those other source systems, onto the specific target system”
“placing the specific source system on the specific target system”	“moving applications and/or data from the specific source system onto the specific target system”
“issue instructions to place the at least one source system on the at least one target system in accordance with the determined placement”	“issue instructions to move applications and/or data from the at least one source system onto the at least one target system in accordance with the determined movement”

a. Densify’s Opening Position

The claim terms “place”/“placed”/“placing”/“placement” should be given their plain and ordinary meaning. The verb “place” is a common dictionary term, defined as “to put in or as if in a particular place or position,” “to assign to a position in a series or category,” Merriam-Webster's Dictionary & Thesaurus at <https://www.merriam-webster.com/dictionary/place> (Ex. B-1), and used in everyday vernacular such that it needs no further construction. VMware’s construction improperly attempts to narrow the “place” terms by replacing them with “move.” However, the patent does not limit the scope of “place” terms to only the word move; the specification makes clear the patentees did not define or limit the scope of these terms during prosecution. *Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 906 (Fed. Cir. 2004) (“Even when the specification describes only a single embodiment, the claims of the patent will not be read restrictively unless the patentee has demonstrated a clear intention to limit the claim scope

using words or expressions of manifest exclusion or restriction.” (internal quotes omitted)); *see also* D.I. 839 at 3 (same).

In the context of the claims, the patents teaches that the “place” claim terms can involve more than the word “move,” and provide several other meanings that fall within the plain and ordinary meaning. For example, the specification describes several examples for the claim terms, such as consolidating, migrating, transferring, moving, combining, stacking, forward consolidating, etc. Below are highlighted examples from the specification:

- “[O]rganizations have become increasingly concerned with such redundancies and how they can best achieve **consolidation** of capacity to reduce operating costs.” ’492 pat. at 2:3-5 (emphasis added); ’459 pat. at 2:5-7.
- “OS-Level application **stacking** involves **moving** the applications and data from one or more systems to the consolidated system.” ’492 pat. at 2:18-20 (emphasis added); ’459 pat. at 2:20-22.
- “The analysis employs a workload **stacking** model **to combine** the source workloads onto the target system. The **combined** workloads are then evaluated using threshold and a scoring algorithm to calculate a compatibility score for each workload type.” ’492 pat. and ’459 pat., each at 17:56-60 (emphasis added).
- “The compatibility is evaluated in two directions, e.g. for a Server A and a Server B, **migrating** A to B is considered as well as **migrating** B to A.” ’492 pat. at 24:63-67 (emphasis added); ’459 pat. at 24:65-67.

See, e.g., ’492 pat. at 6:5-9 (“[H]ypothetical source systems can be . . . “forward consolidated” onto existing target systems.”); ’492 pat. at 13:37-40 (“multiple sources transferred to a single target”); ’492 pat. at 19:6-10; 20:29-31; 20:31-37; 31:66-32:17; 32:52-64; ’459 pat. at 20:27-44; 31:1-10; 31:11-19; 31:42-45; 32:40-45; 33:34-49; 36:59-61; 37:58-62.

The file histories for both the ’492 and ’459 patents confirm the patentee did not limit the scope of the “place” claim terms. *See e.g.*, ’492 pat. File History at DSYCC736590-600, DSYCC7364517-523, DSYCC736467-474, DSYCC736407-416, DSYCC736325-332; ’459 pat.

File History at DSYCC748537-550. In addition, these “place” terms are commonly used within the virtualization industry. Those skilled in the art understand these terms and the context of the claim language, e.g., those skilled in the art would understand that virtualization systems can be consolidated by placing source systems on target systems. *See, e.g.*, ’492 pat. File History at DSYCC736564-587, DSYCC736617-642; DSY01133890-912; DSY00749096-132.

Accordingly, plain and ordinary meaning is appropriate for the “place” claim terms.

b. VMware’s Answering Position

Each term above includes a variation of “placing” a source system onto a target system—but the specification never uses the word “place” at all, let alone to describe source and target systems. The specification, however, is instructive on what the *place* limitations mean; it defines source systems, target systems, and their interactions *only* in the context of moving applications and/or data from a source system to a target system:

[A] source system refers to a system *from which applications and/or data are to be moved*, and a target server or system is a system *to which such applications and/or data are to be moved*. For example, an underutilized environment having two systems 16 can be consolidated to a target system (one of the systems) by *moving applications and/or data from* the source system (the other of the systems) to the target system.

(’492 patent at 5:56-63 (emphasis added).) Accordingly, to the extent the *place* terms find support in the specification, they must be construed in the context of the full limitations and in view of the specification, as VMware proposes.

By contrast, Cirba seeks to construe the word “place” and its derivatives in isolation using general dictionary definitions (e.g., “to put in or as if in a particular place or position; to assign to a position in a series or category”). (*Supra* § II.D.3.a at page 45.) Cirba’s approach encompasses plainly unintended scenarios, such as physically “placing” a physical source computer on top of a physical target computer on a table. *See Phillips v. AWH Corp.*, 415 F.3d

1303, 1322 (Fed. Cir. 2005) (en banc) (warning that “the use of the dictionary may extend patent protection beyond what should properly be afforded by the inventor’s patent”); *Liebscher v. Boothroyd*, 258 F.2d 948, 951 (C.C.P.A. 1958) (“Indiscriminate reliance on definitions found in dictionaries can often produce absurd results.”).

Cirba’s construction has another problem. If a “source system” is a “system *from which* applications and/or data are to be moved,” and a “target system” is a “system *to which* applications and/or data are to be moved,” then what does it mean to “place” a source system on a target system if the source system itself never moves? The specification states only (and Cirba agrees) that the applications and/or data *from within* the source system are to be moved. Cirba acknowledges the contradiction, alternatively and inconsistently arguing that “place” in the context of the virtualization industry should apply. But neither the claims nor specification describing “place” in the context of virtualization or suggest that “placing” includes copying an entire preexisting VM to a host computer.

The specification provides no support for Cirba’s construction of “place.” Cirba’s supposed synonyms from the specification (*i.e.*, “consolidating, migrating, transferring, moving, combining, stacking, forward consolidating, etc.” (*supra* § II.D.3.a at pages 45-46)) are used only in the context of moving *applications and/or data* from a source system to a target system. For example, consistent with the express definition of “source system” (as Cirba acknowledges), the specification describes moving applications and data from a source system, either by transferring them directly onto the target system (“application stacking”) or by creating a new VM on the target system and moving the applications and data *from* the source system to the new VM (“virtualization”). (’492 patent at 2:18-27 (“OS-Level application stacking involves *moving the applications and data from one or more systems to the consolidated system. . .*.”) (emphasis

added)); *id.* at 2:12-13 (“Virtualization involves virtualizing a physical system as a separate guest OS instance on a host machine.”); *id.* at 5:56-59, 6:38-40, 32:52-53.) But the specification never describes moving or transferring a complete, intact source system.

At core, Cirba seeks to import a term from the later-filed ’687 patent—which focused on moving pre-existing VMs as complete entities among host systems—into the ’459 and ’492 patents. These patents do not focus on virtual systems at all.⁸ In the specification of its ’687 patent, which *does* describe moving pre-existing VMs among hosts systems, Cirba used the word “place” and its derivatives extensively. (’687 patent at Fig. 44; *id.* at 3:26-36 (“evaluating the placement of virtual machines in the virtualized environment using the data sets . . . and identifying the existence of virtual machines with suboptimal placements to enable replacement of the virtual machines”); *id.* at 34:15-25 (“Specifically, the analyses can be performed and scheduled to assist in governing, optimizing and planning the placements of virtual machines in the virtual environment.”); *id.* at 2:41-50, 32:38-46; 34:54-35:2; 35:8-36:10, cls. 1-3, 7, 10, 12, 13, 16, 18.) Cirba did not use the term in the ’492 and ’459 patent specifications because, as named inventor Mr. Hillier⁹ admitted in deposition, the ’492 and ’459 patents were not intended to cover virtualization. (*See* Ex. A-14, 356:3-359:24 (explaining that first filing for ’492 and ’459 patents was not virtualization focused).)

When discussing the interactions between source and target systems, the only technique that the specification describes for the “place” terms and their alleged synonyms is the planned

⁸ Cirba first added “placement of source systems onto target systems” to the ’492 application’s pending claims via amendment in October 2018—roughly six months before Cirba sued VMware, and after it had begun contemplating doing so. (D.I. 1032-7 at 80; TT at 665:1-14, 687:14-688:6 (Cirba explaining that August 2018 is when it started considering suing VMware).) The other *place* terms were amended into the claims after Cirba sued VMware. (*See, e.g.*, D.I. 1032-7 at 106 (amending claims to add “placing source systems on target systems” to claim 1).)

⁹ Mr. Hillier is a named inventor on all three of the patents.

movement of applications and data from a source system to a target system. The *place* terms therefore should be construed as VMware proposes, consistent with those disclosures.

c. Densify’s Reply Position

VMware seek to re-define “place” (and its formatives) as “move” (and its formatives) and “source system” as “applications and/or data.” Both disregard the plain claim language, improperly limiting the plain meaning of the terms, and lack support in the intrinsic record.

“Place” is not limited to “move.” The plain meaning of “place” is not “move.” The common definition for “place” (verb) is “to put in a particular place” or “to distribute in an orderly manner.” *See* Ex. B-10 (Merriam-Webster Dictionary (2004)) at 549. “Move” is not synonymous with “place.” *See* Ex. B-11 (Roget’s II The New Thesaurus (3rd ed. 2003)) at 736. The specification does not redefine or restrict the common meaning of “place,” which appears only in the claims. Rather, the specification talks about “*arranging*” systems as well as “moving” them. For example: “The preferred output shown in FIG. 4 arranges systems 16 in the map 32 such that 100% compatibility exists....” ’459 patent, 27:9-12; FIG. 4.; ’492 pat., 26: 47-48; FIG. 4. VMware’s construction (which it will argue excludes instances where source systems are placed onto target systems for the first time, e.g., the formation of a new container) would also exclude specific disclosed embodiments, including scenarios where “hypothetical source systems can be used to simulate the case where a new application is to be introduced into the environment and ‘forward consolidated’ onto existing target systems **16**.” ’492 pat., 6:69; ’459 pat., 6:7-10.

“source systems” are not limited to “applications and/or data.” The ’492 and ’459 patents expressly contemplate—and the claim language clearly recites—placing entire *source systems*, not only what resides within a source system (e.g., applications and/or data), on target systems. VMware’s proposed construction strips away the “system” portion of the claim,

substituting only the contents of that system in its place. Not only does this read out critical claim language, it is directly contrary to the much broader definition of “system” set forth in the record. The patents plainly disclose “compatibilities between systems **16** based on the parameters to determine if efficiencies can be realized by consolidating *either entire systems 16 or aspects or components thereof.*” ’492 pat., 6:61-64 (emphasis added); ’459 pat., 6:62-65 (emphasis added). Consistent with that understanding, the patents explain the word “system” should be understood to encompass entire systems (e.g., “virtual systems”) in addition to components of systems. For example, the patents teach that “systems” can be entire systems (e.g., “computer systems,” “virtual systems,” or hypothetical systems”). ’492 pat., 5:64-6:15; ’459 pat., 5:65-6:16. The patents further teach that the word “systems” also refers to “entities” including “components” of systems (*supra* at § II.D.1.c(ii)). The patents go on to explain, with reference to Fig. 3, that “transfers” are the movement of such systems onto a target system (’492 pat., 6:34-59; ’459 pat., 6:35-60) and that moving a “system” is one way to move “applications and data” (’492 pat., 2:18-27.; ’459 pat., 2:20-29). Such transfers/transferring are referred to as “moving,” “migrating,” “stacking,” etc., which are actions encompassed by the “place” or “placing” terms used in the claims. *Supra* § II.D.3.a at pages 45-46. Therefore, the patents teach that a “system” that is transferred, moved, migrated, etc. can be an entire system (such as a “virtual system”) and is not limited to movement of applications and/or data from within a system (while the system itself is not transferred, moved, migrated, etc.) as VMware proposes. VMware’s proposed construction is presumably intended to eliminate the possibility that “source system” may read on “virtual machines” (a newly-articulated non-infringement theory). *See* Ex. B-12 (2021-06-04 Ltr. D. Van Nort to W. White) at 2. This ignores the disclosures that contemplate the placement of entire systems.

VMware contends that “[t]hese patents do not focus on virtual systems at all” and “the ’492 and ’459 patents were not intended to cover virtualization” (*Supra* § II.D.3.b at page 49), ignoring the patents’ express disclosure of “physical, **virtual** or hypothetical systems.” *See, e.g.*, ’492 pat., 6:16-18; ’459 pat., 6:17-19. And VMware’s characterization of Mr. Hillier’s deposition (which involved *only* the ’687 patent) is incorrect.¹⁰ Far from confirming the (not-yet-issued) ’492 and ’459 patents were not intended to cover virtualization, Mr. Hillier’s testimony merely states the whitepaper of the ’687 patent provisional application “wouldn’t refer to **rebalancing**” *Supra* § II.D.3.b at page 49, Ex. A-14. The “place” terms clearly contemplate placement of entire systems, and not just movement of applications and/or data from within a system. VMware’s proposed constructions should be rejected in favor of the plain and ordinary meaning of each “place” term.

d. VMware’s Sur-Reply Position

Cirba agrees that the express definition of “source system” includes a “system *from which* applications and/or data are to be moved.” (’492 patent at 5:56-57 (emphasis added); Op. at 10 (emphasis added).) It would not make sense to “place” a source on a target under this definition of “source system.” “From which” indicates that part of the source system remains behind, after its applications or data have been moved off. Leaving the source system behind, and moving its contents elsewhere, addresses the concern that “having too many servers can become a burden,” which the patents aim to solve. (’492 patent at 1:59-61.) Placing entire source systems, by contrast, would do nothing to reduce the number of systems in an environment.

¹⁰ Neither the ’492 patent nor the ’459 patent had issued at the time Mr. Hillier gave this deposition.

Cirba now contradicts the express definition of “source systems,” arguing that the patents “expressly . . . contemplate placing *entire source systems*, not only what resides within a source system (e.g., applications and/or data).” (*Supra* § II.D.3.c at page 50 (emphasis added).) But the patentees could not have intended to override this express definition by referring to “plac[ing]” only in the claims, particularly as the word “place” does not appear in the specification at all.

Cirba continues to argue that “placing” encompasses “consolidating,” “migrating,” “transferring,” “combining,” “stacking,” or forward consolidating” a source system. (*Supra* § II.D.3.a at pages 45-46; *supra* § II.D.3.c at pages 50-51.) But none of those actions requires *placing* an entire source system itself. All involve moving the source system’s contents or workloads (*i.e.*, its applications and/or data) off of it, consistent with VMware’s proposed construction. Cirba’s specification excerpts are not to the contrary. For example, Cirba quotes the specifications’ disclosure that “[t]he analysis employs a workload *stacking* model *to combine* the source *workloads* onto the target system.” (*Supra* § II.D.3.a at page 46 (citing ’492 patent at 17:56-60).) This is consistent with the express definition of “source system,” in which the system’s “workloads” (*i.e.*, its applications or data), are moved from the source system(s) and combined onto the target system. The source itself is not moved.

E. Terms for U.S. Patent No. 10,523,492 (“’492 Patent”) Only

**1. “the systems” / “the source systems” / “the target systems”
(’492 patent: claims 1, 12, 23)**

VMware’s Proposed Construction	Densify’s Proposed Construction
Indefinite for lack of a reasonably certain antecedent basis.	Not indefinite and no construction necessary; antecedent basis is reasonably certain.

a. Densify's Opening Position

The asserted claims of the '492 patent make clear that “the systems” refers to “the source systems” and “the target systems” that are being evaluated using at least one rule set to determine whether or not the systems can be placed together on a specific target system. '492 pat. at 38:37-48. The preamble and the claim language provide the antecedent basis for source systems and target systems, reciting “placing source systems on target systems . . . one or more source systems . . . one or more target systems.” *See Digital-Vending Servs. Int'l, LLC v. Univ. of Phoenix, Inc.*, 672 F.3d 1270, 1275 (Fed. Cir. 2012) (presuming that same term carries “the same meaning throughout all of the claims in the absence of any reason to believe otherwise”). The claim language introduces two types of systems, source systems and target systems, and then refers to them both as systems. That does not make the claims indefinite.

The '492 specification describes both the source systems and target systems as “systems” numerous times. '492 pat. at 5:56-63; 6:2-9. For example, it discusses an “Auto Fit Example” in which “[t]he following example demonstrates how a variant of the auto fit algorithm searches for a consolidation solution for a compatibility analysis comprised of 20 *systems* (S1, S2, S3, . . . S20) where 15 systems (S1-15) have been designated to be *source-only* and 5 (S16-20) to be *target-only*.” '492 pat. at 35:1-7 (emphasis added); *see also* Figs. 1-3; 6:41-44, 6:45-47, 6:51-56, 7:32-35, 33:28-36, 33:56-58, 34:27-35, 34:41-47, 34:48-52.

The plain language of the invention permits a person skilled in the art to be reasonably certain of the antecedent basis for “the systems” as it refers to “the source systems” and “the target systems.” Those skilled in the art would understand how the claim terms are being used with the context of the claim language based on the specification. *See Nautilus, Inc. v. Biosig Instruments, Inc.*, 572 U.S. 898, 901 (2016); *see also Energizer Holdings, Inc. v. Int'l Trade Comm'n*, 435 F.3d 1366, 1370-71 (Fed. Cir. 2006)(“When the meaning of the claim would

reasonably be understood by persons of ordinary skill when read in light of the specification, the claim is not subject to invalidity upon departure from the protocol of 'antecedent basis.'").

Here, the claims, viewed in light of the specification, inform those skilled in the art about the scope of the invention with reasonable certainty.

b. VMware's Answering Position

Claims lacking antecedent basis may be held indefinite. *See In re Packard*, 751 F.3d 1307, 1310 (Fed. Cir. May 6, 2014) (affirming indefiniteness rejection based partly on lack of antecedent basis); *accord In re Altenpohl*, 500 F.2d 1151, 1156 (C.C.P.A. 1974) ("Lack of antecedent basis in a claim could render it invalid under 35 USC 112, second paragraph."); MPEP 2173.05(e). Where a limitation lacks express antecedent basis, courts attempt to determine whether a "reasonably ascertainable meaning" exists because, for example, the antecedent basis is implicit. *See Energizer Holdings v. Int'l Trade Comm'n*, 435 F.3d 1366, 1370-71 (Fed. Cir. 2006). If not, the claim is indefinite. *See Slimfold Mfg. Co. v. Kinkead Indus., Inc.*, 810 F.2d 1113, 1116 (Fed. Cir. 1987); *Collabo Innovations, Inc. v. Omnivision Techs., Inc.*, No. 16-197-JFB-SRF, 2017 WL 3670661, at *8-9 (D. Del. Aug. 25, 2017).

Here, "the systems," "the source systems," and "the target systems" each lack antecedent basis and have multiple different potential interpretations. Each interpretation is facially reasonable but imperfect and has different implications for claim scope. The claims thus are indefinite. *See ChefAm., Inc. v. Lamb-Weston, Inc.*, 358 F.3d 1371, 1374 (Fed. Cir. 2004) ("[C]ourts may not redraft claims, whether to make them operable or to sustain their validity.").

(i) "the systems"

"The systems" (highlighted in yellow below) lacks an antecedent and appears twice in the evaluating limitation of the independent claims.

evaluating [1] one or more source systems against [2] other source

systems and against [3] one or more target systems using at least one rule set that evaluates parameters of the systems to determine whether the systems can or can not be placed together on a specific target system . . . ;

The term could reasonably refer to any of the three, differently-colored groups of systems (or some combination of those groups).

The term's first use is in connection with a rule set that evaluates parameters of "the systems." Cirba contends that the antecedent basis for "the systems" is the source systems and the target systems that are being evaluated—*i.e.*, all three groups highlighted above. (*Supra* § II.E.1.a at page 54.) But Cirba's mapping creates a conflict with the second use of "the systems," *i.e.*, "to determine whether the systems can or can not be placed together on a specific target system."

A "target system" is "a system from which applications and/or data are to be moved," which means that a target system cannot be "placed" on another target system. ('492 patent at 5:56-59.) The second use of "systems" thus can refer only to "one or more source systems," "other source systems," or both (but not to "one or more target systems")—these are the only systems that could be "placed together on a specific target system."¹¹ But this contradicts Cirba's view that the first use of "the systems" includes "one or more target systems."

Were the second use of "the systems" limited to "one or more source systems," "other source systems," or both, however, this would create issues for the first use of "the systems." It would mean that the "at least one rule set that evaluates parameters of the systems" would not evaluate parameters of the "one or more target systems." If target system parameters were not

¹¹ Cirba also fails to explain why "the systems" could not plausibly consist of one, and not both, of the two groups of source systems combined with the "one or more target systems."

used to determine whether source systems might be “placed” together on a specific target system, one would not know whether the specific target is suitable for those source systems.

The consequences of these different and imperfect interpretations are significant, as the “parameters” the rule set(s) evaluate depend on which systems comprise “the systems.” If “the systems” includes only some group of source systems, then only rule sets and parameters for that specific group of source systems are needed. If “the systems” includes both source systems and target systems, however, the narrower claim scope requires rule sets and parameters for both source systems and target systems.

(ii) “the source systems”

“[Placing] the source systems” also lacks antecedent basis and reasonably certain scope:

A computer implemented method for placing [1] source systems on target systems, the method comprising:

evaluating [2] one or more source systems against [3] other source systems and against one or more target systems using at least one rule set that evaluates parameters of the systems to determine whether the systems can or can not be placed together on a specific target system . . . ; and

placing the source systems onto the target systems

It is unclear whether the claim is satisfied if (i) the arbitrary “source systems” of the preamble are placed, (ii) only one single source system is “placed,” (iii) every one of the “one or more source systems” is “placed,” (iv) every of the “other source systems” is “placed,” or (v) every source system in the “one or more” and “other” source systems is “placed.” Stated otherwise, it is impossible to determine *which* “source systems” must be placed to infringe.

(iii) “the target systems”

“The target systems” in the “placing” limitation suffers from a similar ambiguity:

A computer implemented method for placing source systems on [1] target systems, the method comprising:

evaluating one or more source systems against other source systems and against [2] one or more target systems using at least one rule set that evaluates parameters of the systems to determine whether the systems can or can not be placed together on [3] a specific target system . . . ; and

placing the source systems onto the target systems

This term is even more difficult to resolve because “the target systems” is plural, but its most reasonable antecedent (“a specific target system”) is singular. It is unclear whether the claim is satisfied if placement occurs (i) on the ambiguous “target systems” in the preamble, (ii) only on the “specific target system” (which is singular, not plural), (iii) on some or all of the “one or more target systems,” or (iv) on some combination of the “specific target system” and the “one or more target systems.” This, too, makes infringement impossible to assess.

c. **Densify’s Reply Position**

A POSA would have readily ascertained the meanings of these terms, in the context of the claim language and specification. *See* Ex. B-13 at ¶¶ 19-35.

(i) **“the systems”**

VMware’s argument boils down to whether the second use of “the systems” (i.e., in the phrase “to determine whether the systems can or can not be placed together on a specific target system,” includes the recited “one or more target systems.”¹² *See Supra* § II.E.1.b(i) at pages 55-57. It does. VMware concedes that the alleged inconsistency between the first and second uses of “the systems” only exists if “the second use of ‘the systems’ [were] limited to ‘one or more

¹² In a footnote, VMware suggests Densify should have explained “why ‘the systems’ could not plausibly consist of one, and not both, of the two groups of source systems combined with the ‘one or more target systems.’” *Supra* § II.E.1.b at page 56 n.11. If VMware thinks that’s a plausible interpretation, it should have argued so (after all, proving indefiniteness by clear and convincing evidence is VMware’s burden—Densify has no burden to prove definiteness). But VMware did not. Had VMware done so, it would directly conflict with VMware’s position that “the systems” cannot include the “one or more target systems.” That VMware did not square its own argument with that supposed interpretation shows the tenuousness of VMware’s indefiniteness theory.

source systems,’ ‘other systems,’ or both” (*i.e.*, to exclude target systems). *See id.* As explained below, use of “the systems” in the phrase “to determine whether the systems can or can not be placed together on a specific target systems” *does* include “the one or more target systems.”

VMware argues that, because the specification defines “target system” as “a system to which . . . applications and/or data are to be moved,”¹³ it means “a target system cannot be ‘placed’ on another target system.” *Supra* § II.E.1.b(i) at page 56. That is true, but it does not render the claims indefinite. Even under the patent’s definition of “target system,” the claims do not require a target system be “placed on another target system.” Instead, the claims relate to determining whether source systems and one or more target system can be “placed together on a specific target system.” The specification describes such examples with respect to the N-to-1 compatibility analysis, which assesses “the compatibility between N source entities and a single target as defined by a transfer set.” ’492 pat., 30:46-48. The specification explains that “[w]hen stacking multiple source entities onto a target, the source entities and targets coexist in the same operating system environment.” *Id.* at 30:23-32. Because this coexistence involves “inherent sharing” of the environment, “the multi-dimensional analysis considers the N-to-1 inter-compatibility between the source entities and the target,” *i.e.*, the “single target” or the “common target” of a transfer set. *See id.* at 6:41-44, 30:23-32, 30:46-48. Thus, the patent teaches, and a POSA would have understood, that determining “whether the systems can or can not be placed together on a specific target system” involves determining whether source systems and a specific target system can be placed together on that target system or in its environment.

¹³ VMware’s brief incorrectly quotes the specification’s definition of *source system* (“a system from which applications and/or data are to be moved” as the meaning of target system. *Compare* ’492 pat., 5:56-59 with *Supra* § II.E.1.b at page 56.

(ii) “the source systems”

VMware argues “it is impossible to determine *which* ‘source systems’ must be placed to infringe” but does not explain why. *Supra* § II.E.1.b(ii) at 57. But the claims are not unclear as to which source systems are “*the* source systems” – they are the only source systems identified in the body of the claim, *i.e.*, the “one or more source system” and “other source systems” that are evaluated against each another.

VMware suggests “source systems” in the preamble are limitations that are distinguishable from the “one or more source systems” and “other source systems” in the body of the claim. But VMware has not shown by clear and convincing evidence that the “source systems” in the preamble is so limiting. The body of the claim recites “evaluating one or more source systems against other sources systems”; those are the systems to which “placing *the source systems* onto the target systems” refers. This is consistent with the description of “multi-dimensional rule-based compatibility analysis,” which teaches “a transfer set can include multiple sources (N) to be transferred to the target, [and] the analysis may evaluate the compatibility of sources amongst each other (N-by-N).” ’492 pat., 29:59-65; 30:7-8 (“An N-by-N intracompatibility analysis evaluates each source system against each of the other source systems.”).

VMware’s hypothetical questions about how many source systems need to be placed to infringe the claims are mere straw men. Placing “the source systems” requires placing the “one or more source systems” and “other source systems” evaluated against each other; how many of them there are is irrelevant.

(iii) “the target systems”

As with “the source systems,” VMware fails to show the phrase “target systems” in the preamble is a limitation that is somehow distinguished from the “one or more target systems”

and the “specific target system” in the body of the claim. VMware also fails to explain why “a specific target system” is allegedly the “most reasonable antecedent” for “the target systems.” *Supra* § II.E.1.b(iii) at page 58. It is not. As VMware recognizes, “the target systems” is plural, but “a specific target system” is not. *See id.* Therefore, the only reasonable antecedent of the plural “the target systems” is not the singular “a specific target system, but the also plural “one or more target systems.”

d. VMware’s Sur-Reply Position

(i) “the systems”

Cirba’s hopelessly confusing Reply confirms that “the systems” is indefinite. Cirba insists that “the systems” that are evaluated to determine if they “can or can not be placed together on a specific target system” *must include* the “one or more target systems.” (*Supra* § II.E.1.c at pages 58-59.) Yet Cirba recognizes that target systems cannot be placed on each other. (*See supra* § II.D.2.a at pages 42-43.)

To reconcile this contradiction, Cirba suggests that “one or more target systems” and “specific target system” are synonymous in the claims. (*Supra* § II.E.1.c at page 59 (“determining ‘whether the systems can or can not be placed together on a specific target system’ involves determining whether source systems and *a specific target system* can be placed together *on that target system . . .*”) (emphasis added).) But it would make no sense to determine whether a specific target system can be “placed” on itself, as Cirba’s argument would allow. (Ex. A-21 ¶ 18.)

Cirba’s expert Reply declaration further highlights the claim’s indefiniteness. Contradicting Cirba’s position, Dr. Madisetti argues that the meaning of “the systems” depends on the specific evaluation being performed. (Ex. B-13 ¶¶ 22-25.) When evaluating “one or more sources” against “other source systems,” Dr. Madisetti posits that “the systems” refers *only* to

source systems, *i.e.*, the “one or more source systems” and “other source systems,” and does not include target systems at all. (*Id.*) When evaluating “one or more source systems” and “one or more target systems” to find possible placements, Dr. Madisetti posits that “the systems” can “only refer to ‘one or more source systems.’” (*Id.* ¶ 24.) When evaluating “whether ‘one or more source systems’ and ‘one or more target systems’ ‘can be placed together on a specific target system’” “using rule sets,” Dr. Madisetti implies that “the systems” includes “one or more target systems.” (*Id.* ¶ 25.) By conceding that “the systems” can refer to any of the three groups of systems VMware identified, he confirms the claims are indefinite. (Ex. A-21 ¶¶ 21-22.)

(ii) “the source systems”

Cirba has clarified its position that infringement requires that *every one* of the “one or more source systems” and “other source systems” be placed on “the target systems.” Restated, Cirba believes that “the source systems” includes both “one or more source systems” and the “other source systems.” (*Supra* § II.E.1.c at page 60.) While this is one interpretation, it is equally plausible that “the source systems” refers only to the “one or more source systems” being evaluated against the “other source systems” and target systems for placement. (Ex. A-21 ¶¶ 23-24.) As Cirba does not dispute that this interpretation is reasonable, the limitation is indefinite.

(iii) “the target systems”

Cirba also has clarified its position that infringement requires that “the source systems” be placed on every one of the “one or more target systems,” but the claim language contradicts its position. The “evaluating” step assesses whether source systems “can or can not” be placed on “a specific target system” (which presumably is from the “one or more target systems”). To then mandate that the source systems be placed on every target system, as Cirba proposes, would be inconsistent with the claim’s recognition that some source systems “can not” be placed on some target systems. (Ex. A-21 ¶¶ 26-27.) The limitation therefore is indefinite.

**2. “at least one of a compatibility score, and a number of transfers”
(’492 patent: claims 3, 14, 25)**

VMware’s Proposed Construction	Densify’s Proposed Construction
“at least one compatibility score and at least one of a number of transfers”	Plain and ordinary meaning. Alternatively: “a compatibility score and/or a number of transfers”

a. Densify’s Opening Position

The dispute centers around whether an additional “at least one of” should be read into the claim. The relevant claim language states that “the placement solution is selected according to one or more criteria, the one or more criteria comprising at least one of a compatibility score, and a number of transfers.” ’492 pat. Claim 3 at 38:54-57. The plain and ordinary meaning of this phrase is clear: there is one or more criteria, and that one or more criteria can be at least one of either a compatibility score, a number of transfers, or both. To the extent helpful to the Court, a clarifying construction would add “and/or” to show that the criteria can be a compatibility score and/or a number of transfers, which is clearly contemplated by the claims.

The specification confirms that the possible criteria used to select the placement solution include a compatibility score, or a number of transfers, or several other options. ’492 pat. at 33:10-28. In fact, the specification describes in the section “Choose Best Candidate Transfer Set,” that “[t]he criteria employed to select the best transfer set 23 varies amongst the algorithms. Possible criteria include the number of transfers, the compatibility score, general compatibility of entities referenced by set and whether the transfer set target is a target-only.” *Id.* at 34:40-47. The specification further describes how each criteria affects “how the compatibility analysis is performed.” *Id.* at 33:37-40. As such, the optimal workload placement solution can select at least one of the possible criteria: a compatibility score, a number of transfers, or both.

VMware’s construction is incorrect because it requires the optimal workload placement solution to select at least both “possible criteria.” The claims and the specification do not *require* the solution to select at least one “compatibility score” *and* at least one of “a number of transfers.” In fact, the specification discusses that “the best candidate consolidation solution can be selected based on some criteria such as the largest reduction of systems with the highest average transfer set scores.” ’492 pat. at 34:63-65. This suggests that the specification contemplated some of the criteria being selected, not requiring all. Nothing in the claims or specification requires that all criteria must be selected as VMware suggests.

b. VMware’s Answering Position

The rules of grammar require that “at least one of a compatibility score, and a number of transfers” is conjunctive—*i.e.*, that its plain meaning is “at least one compatibility score and at least one of a number of transfers.” *SuperGuide Corp. v. DirecTV Enters., Inc.*, 358 F.3d 870, 885-88 (Fed. Cir. 2004). In *SuperGuide*, the Federal Circuit relied on *The Elements of Style* 27 (4th ed. 2000) when holding that the plain meaning of “at least one of,” followed by elements joined by “and” requires at least one element of each type. *Id.* at 885-86; *see also SIMO Holdings, Inc. v. Hong Kong uCloudlink Network Tech. Ltd.*, 983 F.3d 1367, 1377-78 (Fed. Cir. 2021). The claims here use “at least one of” in connection with a set of criteria, just as in *SuperGuide*.

A different meaning would require that the specification “rebut[] the presumption that the [] patentee intended the plain and ordinary meaning of this language.” *SuperGuide*, 983 F.3d at 887. But the specification expressly discloses selecting a workload placement solution based on *both* a compatibility score and a number of transfers.

For example, the specification explains how “a variant of the auto fit algorithm searches for a consolidation solution.” (’492 patent at 35:1-10.) In the described variant, the inputs

require *both* a “minimum allowable compatibility score” and a minimum and maximum number of “sources per target,” which equates to one or more of a number of transfers. (*Id.* at 35:41-46.)

As another example, the specification explains that the consolidation analysis “seeks the consolidation solution that maximizes the *number of transfers* while still fulfilling the *several pre-defined constraints*,” such as constraints related to compatibility scores. (*Id.* at 33:11-14 (emphasis added); *see also id.* at 33:21-26 (“The auto fit input includes the following parameters: transfer type (*e.g.* virtualize or stacking), *minimum allowable overall compatibility score for proposed transfer sets*, *minimum number of source entities to transfer per target*, *maximum number of source entities to transfer per target*, and quick vs. detailed search for the best fit.”) (emphasis added); *id.* at 36:3-4 (“For example, *the criteria are the greatest number of sources and if tied, the highest score.*”) (emphasis added).)

Even if Cirba were correct that the specification does not *require* both of the claimed criteria to be used when selecting the workload placement solution, this would not change the required interpretation. (*Supra* § II.E.2.a at page 63.) As the Federal Circuit has “repeatedly . . . caution[ed],” a court “should not infer that any particular embodiment is included in a claim when there is probative evidence that sufficiently indicates the contrary.” *SIMO*, 983 F.3d at 1378; *id.* at 1379-80 (construing “at least one of” as conjunctive would not ‘contradict the specification’ simply because it would leave out “some alternative embodiments”). Here, the ’492 specification is consistent with the well-known plain meaning of the structure “at least one of . . . and” The claims should be construed consistent with that plain meaning.

c. **Densify’s Reply Position**

VMware’s construction requires an additional, unstated, “at least one of.” *See Supra* § II.E.2.a at page 63. VMware’s argument that the specification allows for “selecting a workload placement solution based on *both* a compatibility score and a number of transfers,” (*supra* §

II.E.2.b at page 64), is contrary to the plain and ordinary meaning of the term, as exemplified by embodiments in the specification that disclose selection of placement solutions without both criteria. VMware's citation to *SuperGuide* is unhelpful. First, here only two criteria are listed, and not a "series of categories of criteria" as in *SuperGuide*. *SuperGuide Corp. v. DirecTV Enterprises, Inc.*, 358 F.3d 870, 885 (Fed. Cir. 2004) ("at least one of each desired criterion; that is, at least one of a desired program start time, a desired program end time, a desired program service and a desired program type"). Second, in *SuperGuide*, the defendant argued that "every disclosed embodiment of a desired criteria list...teaches a conjunctive list that is consistent with the plain meaning of the claim language." *Id.* VMware apparently concedes that not all disclosed embodiments teach a conjunctive list requiring selection of all criteria, citing two examples and then relying on a fallback argument citing *SIMO Holdings*. *Supra* § II.E.2.b at page 65. The Court should reject VMware's proposed construction as inconsistent with the claim language.

d. VMware's Sur-Reply Position

VMware's construction applies the plain meaning of "at least one of . . . and . . ." consistently with *SuperGuide Corp. v. DirecTV Enters., Inc.*, 358 F.3d 870, 885 (Fed. Cir. 2004). To try to overcome this well-established holding, Cirba alleges that *SuperGuide* does not apply to "only two criteria." (Reply at 16.) Not so. *See IPXL Holdings, LLC v. Amazon.com, Inc.*, 333 F. Supp. 2d 513, 525-26 (E.D. Va. 2004) (applying *SuperGuide* to two-element list).

Cirba also contends that applying *SuperGuide* here would exclude embodiments. But as Cirba's own authority recognizes, a claim need not cover all embodiments. *See SIMO Holdings, Inc. v. Hong Kong uCloudlink Network Tech. Ltd.*, 983 F.3d 1378 (Fed. Cir. 2021) (cited in Reply at 16). This is particularly true for the dependent claims here, which were added late in prosecution. *See Helmsderfer v. Bobrick Washroom Equip., Inc.*, 527 F.3d 1379, 1383 (Fed. Cir.

2008) (“[D]ifferent claims are directed to and cover different disclosed embodiments. The patentee chooses the language and accordingly the scope of his claims.”)

**3. “determining at least one optimal combination of applications operable independently on the target system”
(’492 patent: claims 9, 20, 31)**

VMware’s Proposed Construction	Densify’s Proposed Construction
Indefinite	Not indefinite. Plain and ordinary meaning.

a. Densify’s Opening Position

This term is not indefinite and should be given its plain and ordinary meaning. *See Phillips v. AWH Corp.*, 415 F.3d, 1303, 1312-13 (Fed. Cir. 2005). It is comprised of commonly understood words and can be construed with “little more than the application of widely accepted meaning.” *Id.* at 1314; *Bancorp Servs., L.L.C. v. Hartford Life Ins. Co.*, 359 F.3d 1367, 1372 (Fed. Cir. 2004). The meaning “is fairly inferable from the patent,” *id.* at 1373, and “reflects what the specification shows has been invented.” *All Dental Prodx, LLC v. Advantage Dental Prod., Inc.*, 309 F.3d 774, 779 (Fed. Cir. 2002).

The phrase “determining at least one optimal combination of applications operable independently on the target system” means just what it says. It requires “determining at least one optimal combination” of “applications” that are “operable independently on the target system,” i.e., able to be operated independently of each other on the target system. A POSA would readily understand that “*the* target system” refers to the “specific target system” of claim 1, on which it was determined whether systems can or cannot be placed together. *See, e.g.*, claim 1 (“evaluating one or more source systems . . . against one or more target systems . . . to determine whether the systems can or can not be placed together on a specific **target system** . . . and placing the source systems onto the target systems).

b. VMware's Answering Position

A person of ordinary skill in the art would not be able to determine what it means for applications to be “operable independently on the target system.” Cirba effectively confirms this, as it does not endeavor to explain the intended meaning of “operable independently.”¹⁴

The '492 specification never uses the words “independent” or “independently” when discussing “applications,” leaving the POSITA to guess as to what the applications operate independently *of*. One possibility is that the applications must be “operable independently” of each other. Cirba suggests as much in its opening brief (*supra* § II.E.3.a at page 67), but cites nothing from the specification to support that interpretation. Cirba does not even try to explain how this construction is correct. For example, Cirba does not elucidate how a “combination of applications” could be “operable independently” of one another “on the target system,” if the applications also must be “combined” on that same system, as the claims expressly require.

Nothing in the intrinsic evidence helps reconcile these requirements. To the contrary, the prosecution history discusses source systems (the source of the applications that are moved to a target system) as “operable together” on the target system:

[T]he specific limitation of analyzing both compatibility parameters and source system workload data to ultimately determine which of the systems are *operable together on a target system*, is something that was not previously considered. While previous solutions analyzed sizing requirements and constraints for hardware configurations, these solutions have not assessed compatibility nor which of a plurality of systems are *operable together on a target system*.

(D.I. 1032-7 at 35 (Apr. 21, 2017 Applicant Arguments/Remarks Made in an Amendment)

(emphasis added).)

¹⁴ VMware explained why this claim term is indefinite in its July 29, 2021 Invalidity Contentions and during the parties' August 4, 2021 meet and confer about claim construction.

A second possibility is that the applications must be “operable independently” of the source system(s) on which the applications are currently running. For example, one can imagine an application that must be tied to a specific source system (*e.g.*, a printing application that must be run on a source system connected to a printer) and thus depends on it. Other applications, by contrast, might not be similarly constrained and can operate “independently” (*i.e.*, on any system). But the intrinsic evidence also neither supports nor contradicts this reading.

A POSITA cannot ascertain with reasonable certainty what it means for applications to be “operable independently” or what they must operate independently *of*. The uncertain claim scope (and required activity for infringement) renders the claims indefinite. *Nautilus, Inc. v. Biosig Instruments, Inc.*, 134 S. Ct. 2120, 2129 (2014).

c. Densify’s Reply Position

A POSA would have been able to determine what it means for applications to be “operable independently on the target system.” *See* Ex. B-13 at ¶¶ 36-40. The ’492 patent describes techniques for, among other things, determining “compatibilities between systems . . . to determine if efficiencies can be realized by consolidating either entire systems 16 or aspects or components thereof.” ’492 pat., 6:60-64. The analyzed “systems” may include, for example, virtual systems or applications. *See, e.g., id.* at 5:56-6:32. The patent further discloses that when multiple source entities are transferred to the same target, the source entities may be “assumed to be required to conform to the target” when the target system configuration is “rigid” or “concrete,” meaning there “is little flexibility in accommodating individual application requirements.” *Id.* at 30:14-33; 30:40-44. In contrast, when the configuration is “malleable,” the target system may be “generally adaptable in accommodating source entities,” though there may still be “concrete aspects of the target system.” *Id.*

When source entities are “assumed to be required to conform to the target” because the target system is “rigid” or “concrete” (or has “concrete aspects”), a POSA would understand that the applications are required to be “operable independently” on the target system, because the target system has “little flexibility in accommodating individual application requirements.” ’492 pat., 30:14-33; 30:40-44. That is, the applications must be “operable independently on the target system” because the target may not be able to be reconfigured to accommodate each of their individual requirements. *Id.* This is consistent with the specification: when “the target is deemed to be concrete,” “the multi-dimensional analysis considers the N-to-1 inter-compatibility between the source entities and the target as the primary analysis mechanism,” as the “N-to-1 inter-compatibility analysis assesses each source system against the target.” *Id.* at 30:5-7, 14-33, 40-44.

VMware’s argument that a POSA would not understand “what the applications operate independently *of*” misunderstands the ’492 patent. *Supra* § II.E.3.b at page 68. This is not a question of whether applications must be “‘operable independently’ of each other” or “‘operable independently’ of the source system(s) on which the applications are currently running,” as VMware suggests in its two supposed interpretations. *Id.* Rather, the patent teaches applications are “operable independently on the target system” when, for example, they do not require or rely on individual accommodations (e.g., changes to the target system’s configurations) in order to operate on the target system. *See* ’492 pat., 30:14-33; 30:40-44; Ex. B-13 at ¶¶ 36-40.

d. VMware’s Sur-Reply Position

After asserting that “applications operable independently” means “able to be operated independently *of each other*” in its opening brief (*supra* § II.E.3.a at page 67), Cirba now backtracks, stating that “[t]his is *not* a question of whether applications must be ‘operable independently’ of *each other*” (*Supra* § II.E.3.c at page 70 (emphasis added).) Instead,

Cirba newly alleges that applications are somehow “operable independently” *of the target system* because they “do not require or rely on additional accommodations (*e.g.*, changes to the target system’s configurations) . . . to operate on the target system.” (*Id.*) Cirba’s specification excerpt, which does not use the words “operable” or “independent,” sheds no light on what the phrase means. (Ex. A-21 ¶ 34.)

Cirba’s expert offers only *ipse dixit* to explain why “operable independently” could not refer to the other two possibilities VMware identified, *i.e.*, applications “operable independently” of (i) each other, or (ii) the source systems on which the applications are currently running. (Ex. B-13 ¶ 39 (alleging “disagree[ment]” and that a “POSA would not have seen” the issue the same way).) As the phrase is susceptible to three possible meanings, it is indefinite.

**4. “a 1-to-1 compatibility analysis”
(’492 patent: claims 1, 12, 23)**

VMware’s Proposed Construction	Densify’s Proposed Construction
“an analysis that evaluates the compatibility of every system in a collection of systems as a source system against every system in the collection of systems as a target system on a 1-to-1 basis”	“an analysis that evaluates the compatibility of every possible source-target pair combination in a collection of systems on a 1-to-1 basis”

a. Densify’s Opening Position

The patentees act as their own lexicographer for the term “1-to-1 compatibility analysis.” The ’492 patent indicates that “[t]he 1-to-1 compatibility analysis evaluates the compatibility of every possible source-target pair combination in the collection of systems 16 on a 1-to-1 basis.” ’492 pat. at 7:19-21. It states that “[t]he [1-to-1] compatibility analysis evaluates the compatibility of every specified system as source-target pairs on a 1-to-1 basis” and “produces a compatibility score for each system pair so that analyzing a collection of ten (10) systems produces 10x10 scores.” *Id.* at 9:48-52; *see also id.* at 9:11-12 (“The 1-to-1 compatibility

analysis process B evaluates the compatibility of each system pair on a 1-to-1 basis.”). Thus, “1-to-1 compatibility analysis” would have been understood by a POSA as “an analysis that evaluates the compatibility of every possible source-target pair combination in a collection of systems on a 1-to-1 basis.”

b. VMware’s Answering Position

The parties propose similar constructions for “a 1-to-1 compatibility analysis.”

VMware’s construction brings additional clarity to the term by expressly defining the meaning of “every possible source-target pair combination” (which appears in Cirba’s construction) consistently with the specification. Cirba does not explain how the scope differs under its construction. To the extent Cirba disagrees with VMware’s meaning of “every possible source-target pair combination,” the Court should resolve this dispute now.

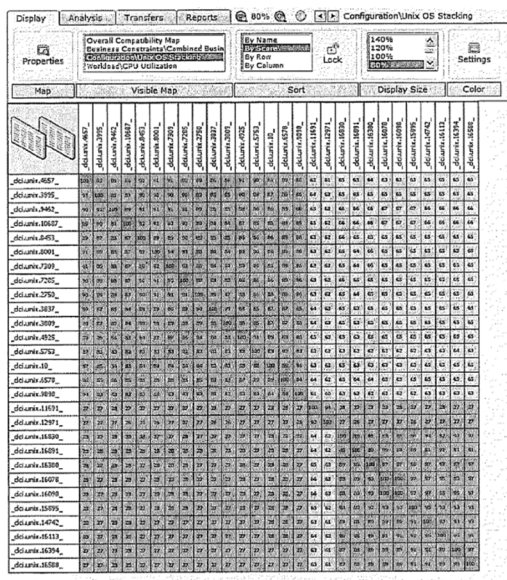


Figure 37

VMware’s construction is consistent with the specification’s disclosure that “[t]he 1-to-1 compatibility analysis evaluates the compatibility of every possible source-target pair combination *in the collection of systems* 16 on a 1-to-1 basis.” (’492 patent at 7:19-21 (emphasis added).) The specification further explains that “the 1-to-1 compatibility analysis of N system[s] computes NxN compatibility scores *by individually considering each system 16 as a consolidation source*

and as a target.” (*Id.* at 26:25-34 (emphasis added).) The scores are “preferably arranged in an NxN map form” (*id.*), such as in the map shown in Figure 37.

This “1-to-1 compatibility map” demonstrates that the 1-to-1 compatibility analysis produces a score for *each* source-target system pairing that exists in a collection of systems—*i.e.*,

every system is listed as a source in the left column and as a target in the top row—as VMware’s construction requires. (*Id.* at 37:17-20, Figs. 39-40, 43.) The specification also depicts the 1-to-1 compatibility analysis as a flow diagram, with dual loops through the collection of systems to generate exhaustive source-target comparisons. (*Id.* at Fig. 17, 16:61-17:3.)

Cirba’s specification citations actually support VMware’s construction. VMware’s construction requires the 1-to-1 compatibility analysis to evaluate the compatibility of every system in a collection of systems as a source system against every system in the collection of systems as a target system, consistent with the specification. (*See, e.g.*, ’492 patent at 9:40-53, 10:1-20, 26:25-34, Figs. 37, 39-40, 43.) Cirba’s citation to column 9, lines 48-51 is fully in accord. (*See supra* § II.E.4.a at page 71 (citing ’492 patent at 9:48-52) (describing “[t]he [1-to-1] compatibility analysis [as] evaluating the compatibility of every specified system as source target pairs on a 1-to-1 basis” and as “produc[ing] a compatibility score for each system pair so that analyzing a collection of ten (10) systems produces 10x10 scores”).)

Cirba’s use of the word “possible” to modify which source-target pairs are evaluated is vague and unhelpful. *See Vivus, Inc. v. Actavis Lab’ys FL Inc.*, No. 14-3786, 2016 WL 3919455, at *3 (D.N.J. July 20, 2016) (“The purpose of claim construction is to reduce ambiguity for the jury.”) For example, it might allow Cirba to circumvent the exhaustive comparisons contemplated by the 1-to-1 compatibility analysis if, for example, analyses before the 1-to-1 compatibility analysis eliminated system pairs from consideration and thus rendered them “impossible.” The specification does not contemplate such pre-filtering for “possible” and “impossible” pairs, instead describing the receipt of a set of systems to analyze and then analyzing them all. (’492 patent at 7:13-25, 9:40-53, 10:5-9, 26:24-28.)

c. Densify’s Reply Position

Densify’s proposed construction is consistent with the patentees’ express definition of this term from the specification. VMware’s proposed construction, cobbled together from disparate portions of the specification, is unhelpful and likely to confuse a jury.

The **1-to-1 compatibility analysis** evaluates the compatibility of every **system in a collection of systems as a source system against every system** in **20** the collection of systems **16 as a target system** on a 1-to-1 basis.

’492 pat., 7:19-21. The red text above shows VMware’s substitutions to the express definition of this term from the specification, but it is unclear how VMware maps its citations to the specification to the substituted text. VMware suggests Densify’s “use of the word ‘possible’ to modify which source-target pairs are evaluated is vague and unhelpful,” but omits the immediately preceding term “every,” which reduces any ambiguity about the contemplated source-target pairs. “[E]very possible source-target pair combination” is how the patentee chose to explain the meaning of this term. *See* ’492 pat., 7:19-21.

d. VMware’s Sur-Reply Position

Rather than explain what “every possible” actually means, Cirba baldly alleges that VMware’s construction is “unhelpful and likely to confuse a jury.” (*Supra* § II.E.4.c at page 74.) But VMware’s construction is even clearer than Cirba’s; it articulates precisely what the patents specified “every possible” to mean. (*See supra* § II.E.4.b at pages 72-73 (citing, *e.g.*, ’492 patent at 9:48-52, 26:25-34, Figs. 17, 37).) Cirba declines to explain why VMware’s clarification of “every possible” is incorrect. To resolve doubts about claim scope, the Court should adopt VMware’s construction.

**5. “an N-to-1 compatibility analysis”
(’492 patent: claims 1, 12, 23)**

VMware’s Proposed Construction	Densify’s Proposed Construction
--------------------------------	---------------------------------

“an analysis that evaluates the compatibility of each of N source systems against a common target system”	“an analysis that evaluates the compatibility of each of N source systems separately and individually against a common target system”
---	---

a. Densify’s Opening Position

The patentees act as their own lexicographer for the term “N-to-1 compatibility analysis,” explaining that “[t]he multi-dimensional compatibility analysis evaluates the compatibility of transfer sets that can involve multiple sources being transferred to a common target.” ’492 patent at 7:26-28. And where “multiple sources (N) [are] to be transferred to the target,” “the [multi-dimensional compatibility] analysis may evaluate the compatibility of . . . each source against the target (N-to-1).” *Id.* at 29:59-65, 30:1-8 (“An N-to-1 intercompatibility analysis assesses each source system against the target.”).

A POSA would have also understood the “N-to-1 compatibility analysis” to separately and individually evaluate each source system against the common target. The specification explains that the multi-dimensional compatibility analysis “evaluates the compatibility of transfer sets that can involve multiple sources being transferred to a common target” and “considers the N-to-1 inter-compatibility between the source entities and the target as the primary analysis mechanism.” *Id.* at 7:26-28, 30:28-32. The patent further explains that “[a] transfer set 23 (see FIG. 3) can be considered one or more transfers that involve a common target, wherein the set specifies one or more source entities[and] the target,” and each “transfer describes the movement of a *single* source entity,” i.e., an individual source entity, “onto a target.” *Id.* at 6:35-44 (emphasis added).

Furthermore, the specification discloses that in the N-to-1 analysis, “compatibility between N source entities and a single target as defined by a transfer set 23” involves “[s]eparately evaluat[ing] each source entity against the target.” *Id.* at 30:46-60 (emphases

added). To illustrate this, the patent describes a hypothetical analysis of “a transfer set t1 [that] comprises of systems s1, s2 and s3 stacked onto s16” that involves “evaluating s1 against s16, s2 against s16 and s3 against s16.” *Id.* at 30:62-67, Table 14. This teaching example shows that each source system (s1, s2, and s3) is evaluated separately (i.e., in a separate evaluation for each source system) and individually (i.e., as a single source system) against the common target system (s16).

Thus, in view of the '492 patent as a whole, a POSA would have understood “N-to-1 compatibility analysis” to mean “an analysis that evaluates the compatibility of each of N source systems separately and individually against a common target system.”

b. VMware's Answering Position

VMware advances the identical construction that Cirba proposed in the '492 patent IPR proceedings. (D.I. 1032-16, Ex. CC at 1719-20.) This construction originates from the specification, which states that, “[a]n N-to-1 intercompatibility analysis assesses each source system against the target.” ('492 patent at 30:5-7.) Cirba's departure from its own PTAB construction to require that evaluating each source system occur “separately and individually” artificially narrows the claims and contradicts the purpose of the N-to-1 compatibility analysis.

The N-to-1 compatibility analysis ensures that multiple source systems can operate *together* on the same target system. (*Id.* at 30:17-19, 30:46-48.) Requiring only separate and individual evaluations in the analysis, as Cirba proposes, would yield information only about specific source-target pairs. This construction would not promote (and indeed would swallow) the N-to-1 compatibility analysis's goal of determining whether moving applications and data from *each* of the N source systems to the *same* target system would be feasible. Separate and individual evaluations could not determine how the N source systems would ultimately operate *together* on the target. Neither the specification nor Cirba explains how one could use such

separate and individual evaluations to arrive at the described output for the N-to-1 compatibility analysis, *i.e.*, a single compatibility score for the transfer set as a whole. (*Id.* at 7:26-31.)

Cirba's argument for including "separately" in the construction (Op. Br. at 18) relies on a single cherry-picked word from the specification describing the "N-to-1 Intercompatibility Score Calculation." ('492 patent at 30:45-61.) But Cirba cannot justify importing one aspect of one step of that score calculation (*i.e.*, "separately evaluat[ing] each source entity against the target") and not all other details (*e.g.*, "determin[ing] whether to count duplicate matched rule items" across different source-target pairs, or "[c]omput[ing] the [N-to-1] score based on the product of all the penalty weights associated with the valid matched rule items" using a specific equation). (*Id.*) Especially as the specification does not support an N-to-1 analysis consisting of purely separate evaluations without considering how to combine the N source systems together, Cirba's selective importation is improper. *See ePlus, Inc. v. Lawson Software, Inc.*, No. 3:09-cv-620, 2010 WL 1779973, at *15 (E.D. Va. Apr. 30, 2010) (rejecting construction based on party "selectively citing portions of the specification helpful to its case while ignoring parts that cut against [its] construction").

Just as Cirba does not explain its insertion of the word "separately," it offers no support for its inclusion of the word "individually." The specification never uses "individually" to describe the N-to-1 compatibility analysis. By contrast, the specification does state that "the 1-to-1 compatibility analysis of N systems computes NxN compatibility scores by *individually* considering each system 16 as a consolidation source and target." ('492 patent at 26:25-28 (emphasis added).) This confirms that, unlike the 1-to-1 compatibility analysis, the N-to-1 is more than a series of separate and individual evaluations. (*Id.* at 29:7-9 ("The multi-dimensional

compatibility analysis extends the original 1-to-1 compatibility analysis that assessed the transfer of a single source entity to a target.”).)

c. Densify’s Reply Position

The dispute here is whether the construction should include “separately and individually.” “Separately and individually” clarifies what the specification means by “evaluates the compatibility of *each* of N source systems *against* a common target.” (emphasis added). The specification teaches the N-to-1 compatibility analysis includes separate evaluations of each individual source system against the common target.’492 pat., 30:2-7. For example, the specification teaches that the N-to-1 compatibility analysis “is performed with respect to a given rule set and involves: 1) Separately evaluate each source entity against the target with the rule set.” *Id.* at 30:46-60. The import of “separately” is shown in the example provided at 30:62-31:15, which describes three separate evaluations of (1) “s1 against S16,” (2) “s2 against s16,” and (3) “s3 against s16” using the rule set. Thus, the patent does not teach, as VMware suggests, evaluating the *combination* of N source systems against the target system, which is why Densify’s construction is appropriate. To use the patent’s example, the patent does not teach evaluating the combination of S1+S2+S3 against S16 using the rule set. Densify’s proposed construction makes that distinction clear, whereas VMware’s ignores it.

Indeed, VMware’s argument that “[t]he N-to-1 compatibility analysis ensures that multiple source systems can operate *together* on the same target system” appears to suggest that the N-to-1 compatibility analysis includes evaluating the compatibility of the combination of S1+S2+S3 against S16 using the rule set in a single analysis—but as explained above, that’s not what the patent teaches. VMware points to the unclaimed “single compatibility score for the transfer set as a whole,” but that score is not the result of evaluating the combination of S1+S2+S3 against S16—it reflects a compatibility score assigned to the transfer set as a whole

based on separate (and individual) evaluations of each source system against the target. '492 pat., 30:45-31:50. But even that unclaimed score is determined by “evaluat[ing] the compatibility of . . . each source against the target (N-to-1),” (*id.* at 29:59-65), which, the patent explains, means evaluating the compatibility of each of N source systems separately and individually against a common target system (*see id.* at 30:45-31:49). Densify’s proposed construction makes that fundamental aspect of the N-to-1 compatibility analysis clear; VMware’s seeks to obfuscate it.

VMware argues the N-to-1 compatibility analysis does not evaluate each source system “individually” against the target system, because systems are “individually” considered in the 1-to-1 compatibility analysis. *Supra* § II.E.5.b at pages 77-78. This is wrong. In a N-to-1 compatibility analysis, each of N source systems is individually evaluated against the target. '492 pat., 30:62-31:15 (evaluating each of S1, S2, and S3 as an individual source system against the target system S16). That a compatibility score is generated based on the individual evaluations does not change the fact each source system is individually evaluated against the target in an N-to-1 compatibility analysis.

d. VMware’s Sur-Reply Position

Cirba’s added requirement that the “N-to-1 compatibility analysis” perform “separate[] and individual[]” evaluations does not clarify the phrase’s meaning, but selectively imports limitations and artificially narrows the claims. Cirba still does not explain why the analysis must include only the first step of the exemplary “N-to-1 Intercompatibility Score Calculation.” ('492 patent at 30:45-61.) As VMware has explained, the exemplary calculation includes the steps of “determin[ing] whether to count duplicate matched rule items” across different source-target pairs and “[c]omput[ing] the [N-to-1] score based on the product of all the penalty weights associated with the valid matched rule items” as part of calculating a single N-to-1 compatibility

score.¹⁵ (*Id.* at 30:52-57.). Cirba also does not explain how “separate” evaluations could determine if source systems can operate *together* on the same target, without further analysis.

F. Terms for U.S. Patent No. 10,951,459 (“’459 Patent”) Only

**1. “the source systems”
(’459 patent: claims 1, 31, 63)**

VMware’s Proposed Construction	Densify’s Proposed Construction
Indefinite for lack of a reasonably certain antecedent basis.	Not indefinite. Plain and ordinary meaning. Alternatively: “the specific source system and those other source systems”

a. Densify’s Opening Position

Read in the proper context, a person of ordinary skill in the art would reasonably be certain that “the source systems” in claims 1 and 31¹⁶ would refer back to “the plurality of source systems” recited earlier in the same claim limitation (see claim 1(b)(2) below).

evaluate compatibility between the specific source system from *the plurality of source systems* and one or more other source systems either already placed on the specific target system, or being evaluated for placement onto the specific target system, to determine if the specific source system can be placed with those other source systems on the specific target system, by evaluating one or more rules that operate against attributes or data relating to *the source systems*[.]

’459 pat. at 40:18-27 (emphasis added); 42:14-35; *Phillips*, 415 F.3d at 1312. Those terms are then presumed to carry the same meaning throughout each claim. *Digital-Vending Servs.*, 672 F.3d at 1275 (presuming that same term carries same meaning in different portions of a claim and across different claims).

¹⁵ Cirba argues that the “single compatibility score” is “unclaimed.” (*Supra* § II.E.5.c at page 78.) It is just as unclaimed as evaluating “separately and individually,” which Cirba seeks to import selectively.

¹⁶ Claim 63 does not include the term “the source systems,” making the basis for VMware’s argument unclear.

The '459 patent describes an example “consisting of *systems* S1 to S20, the compatibility of the *source systems* S5 to S20 can be individually assessed against the transfer set T1. Similarly, the compatibility of the *source systems* S1 to S4 and S11 to S20 can be assessed with respect to the transfer set T2.” ’459 pat. at 30:21-26 (emphasis added). Another example “demonstrates how a variant of the auto fit algorithm searches for a consolidation solution for a compatibility analysis comprised of 20 *systems* (S1, S2, S3, . . . S20) where 15 *systems* (S1-15) have been designated to be *source-only* and 5 (S16-20) to be *target-only*.” *Id.* at 36:20-25 (emphasis added). In other words, the patent teaches, for example, a consolidation analysis involves multiple systems, and a subset of those systems are designated to be source systems and to be target systems.

The language used to describe the source systems and target systems aspect of the invention is parallel in the claims at issue: “evaluate compatibility between the specific source system from the plurality of source systems and one or more other source systems . . . to determine if the specific source system can be placed with those other source systems on the specific target system, by evaluating one or more rules . . . the source systems.” ’459 pat. at 40:18-27. In the context of the patent, it is reasonably certain that “the source systems” in the claims, like the “source systems” in the description of the patent, refer to the plurality of source systems that is designated from the group of systems for the consolidation analysis and solution. *Nautilus*, 572 U.S. at 901. Thus, contrary to VMware’s argument, “the source systems” is not indefinite and should be given its plain and ordinary meaning.

b. VMware's Answering Position¹⁷

Cirba does not dispute that “the source systems” (highlighted below) has no antecedent basis. (*Supra* § II.F.1.a at pages 80-81.) It might reasonably refer to three different groupings or any combination of them: (i) **the specific source system** and at least one of the other source systems; (ii) **the plurality of source systems**; and (iii) **one or more other source systems**:

evaluate compatibility between [1] **the specific source system** from [2] **the plurality of source systems** and [3] **one or more other source systems** either already placed on the specific target system, or being evaluated for placement onto the specific target system, to determine if **the specific source system** can be placed with **those other source systems** on the specific target system, by evaluating one or more rules that operate against attributes or data relating to the source systems;

Cirba's brief illustrates why “the source systems” has multiple reasonable meanings and thus is indefinite. Cirba's alternative construction is “**the specific source system** and **those other source systems**” (options [1] and [3]). Yet it begins by arguing that “the source systems” refers to “**the plurality of source systems**” (option [2]), contradicting itself. (*Supra* § II.F.1.a at page 80.) Moreover, “**the specific source system** and **those other source systems**” does not have the same scope as “**the plurality of source systems**”; the latter may or may not include “**those other source systems**” (*i.e.*, the claim does expressly link the two together).

If “the source systems” instead referred to “**the specific source system** and **those [one or more] other source systems**,” the claim would require each rule to operate against data and attributes of two elements: [1] **the specific source system** and [3] **the other source systems**. That is, a rule that operated against the attributes and data of the “**specific source system**” and one (or all) of “**those other source systems**” might satisfy the claim, but a rule operating only on the

¹⁷ Since serving its brief, Cirba has agreed claim 63 should read “the source systems,” consistent with claims 1 and 31, rather than “the two or more source systems.” (*See supra* § II.F.1.a at page 80 n.16.)

attributes and data of one (or all) of “those other source systems” (and not the specific source system) would not. “The plurality of source systems,” on the other hand, might require only that each rule operate against the data and attributes of one of any of the plurality of source systems, and not against any particular source system, such as the specific source system.

“The systems” also could just as reasonably refer solely to the “one or more other source systems,” excluding “the specific source system.” The rationale for this reading would be that “those other source systems” must provide the antecedent basis for “the systems,” as the two phrases are closest to each other in the claims. Although Cirba presumably disputes this reading, it does not explain why it is incorrect.

The specification does not clarify the appropriate scope of “the source systems.” For instance, the compatibility analysis and auto fit examples that Cirba analyzes (*supra* § II.F.1.a at page 81) do not identify which system is the “specific source system,” much less whether that system is included in “the source systems.” In fact, Cirba’s examples do not confirm whether any of the numbered source systems are compared with other source systems at all, and if so, which source systems’ attributes and data the rule sets operate against. As the claims are susceptible to at least three, equally plausible interpretations of differing scope and nothing in the claims or specification explains which is correct, the claims are indefinite.

c. Densify’s Reply Position

The claims are not unclear as to which source systems are “the source systems” – they are “the specific source system from the plurality of source systems” and the “one or more other source systems” that the claims require be evaluated. This reading of the plain language is further underscored by the claims’ use of the words “evaluated compatibility *between*,” which makes clear that two things are evaluated. This understanding of the claim language is further supported by the specification. *See* ’459 pat., 7:26-28. The Court should reject VMware’s proposed

construction, which is inconsistent with the claim language, and either apply plain and ordinary meaning or adopt Densify’s alternative construction. *See* Ex. B-13 at ¶¶ 41-45.

d. VMware’s Sur-Reply Position

Although Cirba insists that the meaning of “the source systems” is clear, Cirba offered one construction in its opening brief (*supra* § II.F.1.a at page 80 (“the plurality of source systems”)) and another in its Reply (*supra* § II.F.1.c at page 83 (“the specific source system” and “one or more other source systems”).) Cirba still has not explained why “the source systems” could not plausibly refer only to the “those other source systems”—a third option that VMware previously identified (*Supra* § II.F.1.b at page 83).

Claim scope necessarily turns on the meaning, as the claims require “evaluating one or more rules that operate against attributes or data relating to *the source systems*.” Absent clarity as to which “source systems” are at issue, the POSITA would not know which rule(s), attributes, or data are relevant to infringement. (Ex. A-21 ¶¶ 37-40.) The limitation thus is indefinite.

III. DISPUTED CONSTRUCTIONS FOR VMWARE’S PATENTS

A. U.S. Patent No. 9,379,995 (“’995 Patent”)

1. “snapshot” (’995 Patent: claims 1, 9, and 17)

VMware’s Proposed Construction	Densify’s Proposed Construction
“data at a particular moment in time”	“data that contains configuration and resource usage information of a distributed computer system at a particular moment in time”

a. VMware’s Opening Position

The ’995 patent (D.I. 1032-4, Ex. K) discloses inventions related to resource allocation in a distributed computer system, such as a virtual environment. It describes methods that obtain and analyze configuration and usage data from resource-consuming clients in the distributed

computer system, such as virtual machines, and make allocation recommendations for the available resources (*e.g.*, CPU and memory). (*Id.*, 6:10-7:52.) More specifically, the claimed inventions “generate a resource allocation recommendation” by “iteratively traversing a resource hierarchy in the distributed computer system.” (*Id.*, claims 1, 9, 17.) As discussed below, resource hierarchies are resources (*e.g.*, CPU and memory) that are organized in ranks, layers, or a tree structure. (*See also id.*, FIG. 4, 9:50-10:8.)

The parties agree that a “snapshot” is data “at a particular moment in time.” But the remainder of Cirba’s proposal renders other claim language superfluous. The claim step at issue is: “obtaining a target resource allocation and *a snapshot of the distributed computer system, wherein the snapshot includes configurations and resource usage information of at least some components of the distributed computer system.*” (’995 patent, claims 1, 9, 17.) As the italicized portion shows, Cirba’s additional proposed language already appears in the claims. Under Cirba’s construction, the italicized language would become superfluous. *See IpLearn, LLC v. Kenexa Corp.*, No. 11-825-RGA, 2013 WL 5730610, at *4 (D. Del. Oct. 22, 2013) (rejecting proposed limitation that would make subsequent claim limitations superfluous); *see also Douglas Dynamics, LLC v. Buyers Prods. Co.*, 717 F.3d 1336, 1350 (Fed. Cir. 2013) (superfluous construction disfavored).

b. Densify’s Answering Position

Densify’s proposed construction of “snapshot” reflects the patentee’s express definition: “As used herein a snapshot of a distributed computer system contains at least configuration and resource usage information of the distributed computer system at a particular moment in time.” ’995 patent at 7:53-56.¹⁸ “That definition controls.” *Martek Biosciences Corp. v. Nutrinova*,

¹⁸ Courts routinely find that the phrase “as used herein” signals an express definition. *See Chevron U.S.A. Inc. v. Univ. of Wyoming Rsch. Corp.*, 978 F.3d 1361, 1366 (Fed. Cir. 2020); *Housey*

Inc., 579 F.3d 1363, 1380 (Fed. Cir. 2009) (collecting cases); *see also Fonar Corp. v. Johnson & Johnson*, 821 F.2d 627, 632 (Fed. Cir. 1987) (inventors “must use words in the same way in the claims and in the specification.”). The claims’ “wherein” clause is not rendered superfluous because the terms operate as a single limitation, mutually reinforcing that the claimed “snapshot” must contain *at least* “configuration[(s)] and resource utilization information,” as distinguished from other types of data.¹⁹ *See* ’995 patent at 7:53–8:14 (describing other data types); *Bell & Howell Document Mgmt. Prod. Co. v. Altek Sys.*, 132 F.3d 701, 707 (Fed. Cir. 1997).

VMware’s argument that Densify’s construction renders other claim language “superfluous” because it contains “language [that] already appears in the claims” fails for multiple reasons. First, Densify’s “snapshot” construction does not render the subsequent “wherein” clause superfluous because the “wherein” clause requires that the “configuration and resource information” are of “*at least some components of*” the distributed computer system, which is not included in Densify’s “snapshot” construction.

Second, the maxim that claims generally should not be interpreted to render their terms superfluous is not violated where, as here, the allegedly superfluous terms are part of the same limitation of the same claim. *See World Class Tech. Corp. v. Ormco Corp.*, 769 F.3d 1120, 1125

Pharms., Inc. v. Astrazeneca UK Ltd., 366 F.3d 1348, 1354 (Fed. Cir. 2004); *Bristol-Myers Squibb Co. v. Aurobindo Pharma USA Inc.*, No. CV 17-374-LPS, 2018 WL 5077895, at *4 (D. Del. Oct. 18, 2018); *Andover Healthcare, Inc. v. 3M Co.*, No. CV 13-843-LPS, 2015 WL 2227786, at * 5 (D. Del. May 11, 2015).

¹⁹ VMware’s U.S. Patent No. 10,069,752 (the “’752 patent”) discloses *the same express definition* of the term “snapshot” as that found in the ’995 patent (D.I. 737-3, ’752 patent at 7:39-42), and Claim 1 of the ’752 patent includes the same modifying language for “snapshot”: “performing remote resource allocation analyses on the computer system snapshots, each of the computer system *snapshots including configurations and resource usage information of at least some computer system components . . .*” (*id.* at 15:57-61) (emphasis added). For the ’752 patent, the parties agreed to the construction Densify proposes here (D.I. 788 at 2), but VMware refuses here, despite the fact that the terms, the relevant surrounding claim language, and the supporting disclosures are identical.

(Fed. Cir. 2014) (“The doctrine of claim differentiation creates a presumption that distinct claims, particularly an independent claim and its dependent claim, have different scopes.”) (citing *Kraft Foods, Inc. v. Int’l Trading Co.*, 203 F.3d 1362, 1368 (Fed.Cir.2000)); *Bell & Howell*, 132 F.3d at 705–707 (rejecting argument that a construction “would render the [term] superfluous because [part of the construction] is already recited in the claims”); *see also Arendi S.A.R.L. v. LG Elecs., Inc.*, 2019 WL 3891150, at *13 (D. Del. Aug. 8, 2019) (J. Stark) (noting that differentiation between “an independent claim and one of its dependent claims [is] a situation in which claim differentiation is at its strongest”). In *Bell & Howell*, the limitation at issue was “integrally bonded . . . free of adhesives.” The district court concluded that “‘integrally bonded’ could not be construed to mean bonded without adhesive, because this would render the claim language ‘free of adhesive’ superfluous.” *Id.* at 703–704. The Federal Circuit disagreed, finding that the intrinsic evidence made clear that being “integrally bonded” meant that the bond was “free of adhesive.” *Id.* at 706–707. The Federal Circuit found that “‘integrally bonded . . . free of adhesive’ therefore operated as a single limitation such that “[b]eing ‘integrally bonded’ and ‘free of adhesive’ are mutually reinforcing definitions rather than being superfluous.” *Id.* at 707.

This case is similar. The term “snapshot” appears in the same limitation as the allegedly superfluous language: “obtaining . . . a snapshot of the distributed computer system, wherein the snapshot includes configurations and resource usage information of at least some components of the distributed computer system.” ’995 patent, claim 1. The specification defines “snapshots” as “contain[ing] at least configuration and resource usage information of the distributed computer system” (’995 patent at 7:53-56) making clear that the phrase operates as a single limitation that mutually enforces what the snapshot data is. Under *Bell & Howell*, that is not superfluous.

VMware’s citation to the dissent in *Douglas Dynamics, LLC v. Buyers Prods. Co.* for the general principle that “superfluous constructions [are] disfavored” is inapposite. *Supra* § III.A.1.a at page 85 (citing 717 F.3d 1336, 1350 (Fed. Cir. 2013)). The majority in *Douglas Dynamics* found the allegedly superfluous construction was *not* superfluous. 717 F.3d at 1343. VMware’s citation to *IpLearn, LLC v. Kenexa Corp.* fares no better. *Supra* § III.A.1.a at page 85 (citing No. 11-825-RGA, 2013 WL 5730610, at *4 (D. Del. Oct. 22, 2013)). In *IpLearn*, the Court found a construction seeking to import a different limitation from a dependent claim into an independent claim would render the limitation superfluous in the dependent claim. 2013 WL 5730610, at *4 n.3. *IpLearn* is inapposite to this case, which involves mutually reinforcing terms that are part of the same limitation of the same claim. *See Bell & Howell*, 132 F.3d at 707.

c. VMware’s Reply Position

Cirba does not dispute that the parties’ constructions have identical scope. Yet it insists on adding superfluous language—“that contains configuration and resource usage information of a distributed computer system”—even though nearly identical language already appears in the claims. That the parties agreed to Cirba’s construction of “snapshot” for a *different* patent is irrelevant to whether the additional language makes sense here. (*See supra* § III.A.1.b at page 86 n.19.)

Bell & Howell Document Management Products Co. v. Altek Systems, 132 F.3d 701 (Fed. Cir. 1997), does not help Cirba. There, the Federal Circuit found that the two different claim terms, “integrally bonded” and “free of adhesive,” were “mutually reinforcing” rather than “being superfluous.” *Id.* at 707 (construing “integrally bonded . . . free of adhesive” as a whole, as “the ribs bond to the panels without the use of a separate adhesive layer between the ribs and the panels”). No similar justification for including superfluous language exists here.

d. Densify’s Sur-Reply Position

VMware cannot deny that the specification provides an express definition of “snapshot” that matches Densify’s construction. That definition controls. *Martek Biosciences Corp. v. Nutrinova, Inc.*, 579 F.3d 1363, 1380 (Fed. Cir. 2009) (collecting cases); *see also Fonar Corp. v. Johnson & Johnson*, 821 F.2d 627, 632 (Fed. Cir. 1987) (inventors “must use words in the same way in the claims and in the specification.”). VMware’s argument based on the dissent in *Douglas Dynamics, LLC v. Buyers Prods. Co. Supra* § III.A.1.a at page 85 (citing 717 F.3d 1336, 1350 (Fed. Cir. 2013)) should be rejected here for the same reasons as cited by the majority in that case.

2. “iteratively traversing a resource hierarchy” (’995 Patent: claims 1, 9, and 17)

VMware’s Proposed Construction	Densify’s Proposed Construction
<p>Not indefinite. This term should not be construed as it takes its plain and ordinary meaning.</p> <p>Alternatively, this term should be construed as “repeatedly moving along resources that are organized in a hierarchy, such as in ranks, layers, or a tree structure.”</p>	<p>Indefinite for failure to point out with particularity and distinctly claim the subject matter such that one of ordinary skill in the art would be reasonably apprised of the bounds of the asserted claims.</p>

a. VMware’s Opening Position

This phrase requires no construction, as its words take their plain and ordinary meaning and the phrase is not indefinite. Alternatively, if the Court construes this phrase, it should be construed to reflect the plain and ordinary meaning: “repeatedly moving along resources that are organized in a hierarchy, such as in ranks, layers, or a tree structure.”

“Resource Hierarchy.” The phrase “resource hierarchy” has a plain and ordinary meaning in the art. A “resource” is “[a]ny part of a computer system or a network, such as a disk

drive, printer, or memory, that can be allotted to a program or a process while it is running.” (Ex. A-1 (*Microsoft Computer Dictionary* (5th ed. 2002)) at 451; *see also* Ex. A-2 (*Oxford Dictionary of Computing* (6th ed. 2008)) at 435 (defining “resource” as “[a]ny of the component parts of a computer system and the facilities that it offers,” stating “[a]ll computer systems” have “one or more processors” and “some form of memory”).) “Hierarchy” means “a group of things arranged in order of rank” (Ex. A-3 (*Newton’s Telecom Dictionary* (25th Anniversary ed. 2009)) at 554) or “a way in which objects or data or structures are organised, usually with the most important or highest priority or most general item at the top, then working down a tree structure” (Ex. A-4 (*Dictionary of Computing* (6th ed. 2010)) at 160). The plain meaning of “resource hierarchy” thus is resources (e.g., CPU and memory) organized in ranks, layers, or a tree structure.

The specification supports this plain meaning. It describes a “resource pool” embodiment of a resource hierarchy as “aggregated resources of the distributed computer system [] partitioned in a hierarchical fashion using a resource pool (RP) abstraction.” (’995 patent, 9:50-52.)²⁰ Such resources include CPU and memory. (*Id.*, 6:10-36.) Partitioning resources hierarchically “allows service providers (e.g., cloud providers) to group related clients (e.g., VMs) and provide performance isolation among resource groups.” (*Id.*, 9:52-55.)

The ’995 patent further describes a specific example of the “resource pool” embodiment in FIG. 4, which depicts “a hierarchical organization of a virtualized computing environment.”

²⁰ “A resource pool is a logical abstraction for flexible management of resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources.” (Ex. A-5 at Ex. B at 12; *see also id.* at 45-51.) “A resource pool can contain child resource pools, virtual machines, or both. You can create a hierarchy of shared resources. Resource pools and virtual machines that are at the same level are called siblings.” (*Id.* at 45.)

(*Id.*, FIG. 4, 9:60-10:8.) This exemplary resource hierarchy includes a resource cluster at the root (highest) level and organized layers of resource pools and client virtual machines within the cluster. (*Id.*, FIG. 4.) In the FIG. 4 example, the resource pools are organized into two different layers (the first and second layers) within the cluster, and above the client VM layer (the third layer). The resources (*e.g.*, CPU and memory) available in a parent resource pool are shared by that parent's children (*e.g.*, child resource pools or client virtual machines):

The resource cluster includes hierarchical organized RPs 402, 404, 406, 408 and VMs 420. In particular, the resources available in the resource cluster are divided into RPs 402, 404, which are located in a first layer that is directly underneath the resource cluster 400. The resources in RP 402 are shared by VMs 420-1, 420-2, which are located in a second layer that is underneath the first layer. The resources in RP 404 are shared by RPs 406, 408, which are located in the second layer. The resources in the RP 406 are shared by VMs 420-3, 420-4, 420-5, which are located in a third layer that is underneath the second layer. The resources in RP 408 are shared by VMs 420-6, 420-7, which are located in the third layer.

(*Id.*, 9:62-10:8; *see also id.*, FIG. 4.)

“Iteratively Traversing a Resource Hierarchy.” “Iteratively traversing” similarly takes its ordinary meaning. The word “iteratively” is an ordinary English word meaning repeatedly—*i.e.*, more than once. (*See* Ex. A-6 (*Merriam-Webster’s Collegiate Dictionary* (11th ed. 2007)) at 666) (defining “iterative” as “involving repetition . . . — iteratively *adv*”).) The word “traverse” means “move or pass along or through.” (*Id.* at 1332.) Thus, by its plain meaning, the term “iteratively traversing a resource hierarchy” means repeatedly moving along resources that are organized in a hierarchy, such as in ranks, layers, or a tree structure.

The specification illustrates that “iteratively traversing a resource hierarchy” takes this plain meaning. In one embodiment, the disclosed method iteratively traverses the resource hierarchy by starting from the bottom layer of the hierarchy and “traversing higher” in the hierarchy. (’995 patent, 13:6-10, 15:28-31.) This traversal moves along the hierarchy of resource pools within the cluster more than once, *e.g.*, by moving from a client VM to the “VM’s

parent RP node and other ancestor nodes higher up in the RP tree as necessary.” (*Id.*, 13:1-30.)

As the patent illustrates using FIG. 4, an exemplary method starts from a virtual machine (420-3) at the third layer, moves to its parent resource pool (406) in the second layer, and then moves to a sibling resource pool (408) in the second layer. (*Id.*, 15:31-44.)

The dependent claims also demonstrate that either no construction is needed or VMware’s alternative construction is correct. Consistent with the example discussed above (*id.*), claims 2, 10, and 18 limit “iteratively traversing a resource hierarchy” by requiring analysis of an entity in the bottom layer of the hierarchy and then subsequent analysis of each parent resource of the entity. (’995 patent, claims 2, 10, and 18.) In that case, the traversal starts from the bottom layer. By contrast, broader claim 1 (from which claim 2 depends) also includes, within its scope, “iteratively traversing” starting from a higher layer in the resource hierarchy and moving down, up, or within a layer of the resource hierarchy.

b. Densify’s Answering Position

The ’995 patent does not provide “reasonable certainty” as to what “iteratively traversing a resource hierarchy” means because, even when considered in view of the specification, the phrase “might mean several different things and no informed and confident choice is available among the contending definitions.” *Nautilus, Inc. v. Biosig Instruments, Inc.*, 572 U.S. 898, 910 (2014) (for purpose of § 112, ¶ 2, “a patent’s claims [are] viewed in light of the specification and prosecution history”); *Media Rts. Techs., Inc. v. Cap. One Fin. Corp.*, 800 F.3d 1366, 1371 (Fed. Cir. 2015) (citing *Nautilus*, 572 U.S. at 911 n.8).

As a first matter, the parties agree that “resource hierarchy” means “resources organized in ranks, layers, or a tree structure.” *Supra* § III.A.2.a at page 90. The issue then narrows to what the patent means by “iteratively traversing” a resource hierarchy. As detailed below, the patent teaches *three* possible meanings: one, where “iteratively traversing” means simply

“repeatedly moving” along the hierarchy, *i.e.*, making more than one move among nodes (*e.g.*, node-to-node-to-node); two, where “iteratively traversing” occurs as long as the adjustment operations are performed at more than one node, *e.g.*, making adjustments at two nodes (which requires only a single movement from a first node to a second node); and three, where “iteratively traversing” requires both repeatedly moving along the hierarchy *and* performing the adjustment operation at each node. A person of ordinary skill in the art (hereinafter “POSA”) would not be able to ascertain with reasonable certainty which meaning controls, thereby rendering the term indefinite.

VMware argues that “iteratively traversing” should be given its “plain and ordinary” definition and proposes (as an alternative) a construction for what that is: “repeatedly moving....” VMware cites to disclosures that describe a “resource allocation diagnosis algorithm travers[ing] the resource hierarchy, starting from the bottom VM layer.” ’995 patent at 13:6–10; Fig. 4 (showing a “root” and three layers). VMware additionally identifies disclosures that teach “traversing higher in the RP tree of a resource cluster,” adjusting parameters of a VM’s “parent RP 406 and/or its next higher ancestor RP 404,” and “travers[ing] higher up in the RP tree to the sibling RP 408 of the RP 406.” *Id.* at 15:28–36. VMware argues that per these disclosures, a POSA would understand “iteratively traversing” to mean “repeatedly moving” along the hierarchy. Under this construction, the term could only be infringed if there are “repeated” movements on an RP tree. *Id.* at 15:28–36, 13:6–10, Fig. 4. Significantly, a *single* movement would not qualify – the movements must occur more than once to be “repeated.”

But VMware ignores other disclosures in the patent that teach a second possible meaning for “iteratively traversing,” wherein a single movement *could* be sufficient to meet the claim language. Specifically, the ’995 patent specification provides examples where “traversal”

involves carrying out an *adjustment* sequence on a first node and then iterating the sequence on a second node (*i.e.*, at a next higher level in a hierarchy of resources), as opposed to “repeatedly moving” levels. Ex. B-2 at ¶¶ 23-26. Specifically, the ’995 patent defines a resource allocation diagnosis algorithm as including three ordered procedures: “AdjustLimit,” “AdjustShares,” and “AdjustReservations.” *Id.* at 13:13–41. The “limit is adjusted first,” for a current node, and “[i]f [it] cannot be adjusted to a required level, other resources such as the shares and the reservation are adjusted.” *Id.* The ’995 patent teaches that “[d]epending on the RP tree hierarchy, these procedures *can be also* applied to the target VM’s parent RP node and other ancestor nodes higher up in the RP tree,” and that the algorithm “iterates” (*i.e.*, repeats) the sequence of operations, “for all the nodes in a resource cluster.” *Id.* (emphasis added); Ex. B-2 at ¶ 26. “For each parent of the current node in the resource pool hierarchy, the limit is adjusted first,” and “[i]f [it] cannot be adjusted to a required level, other resources such as the shares and the reservation are adjusted.” *Id.* The “depending on the RP tree hierarchy” and “can also be applied” language would indicate to a POSA that “iteratively traversing a resource hierarchy” need not require multiple or repeated “movements” as VMware posits — rather, the traversal of one layer (or node) higher in the hierarchy that involves repeating the adjustment process (*i.e.*, on a first and a second node) may be enough.

The critical difference comes in the “iteratively” requirement of the claim: iteratively “traversing” under the second possible meaning can occur even where there is only a single move from one level to the next. Essentially, it is possible to have *one* movement between two layers that involves *two* adjustments (adjustment at the first level, move up, adjustment at the second level). Under VMware’s possible construction, this scenario would not meet the letter of the claim, because more than one movement is required to “iteratively” traverse; but under the

other possible construction, it could, because “iteratively” traversing occurs through the multiple adjustments, and the number of movements is irrelevant. Ex. B-2 at ¶ 27

VMware also ignores that the patent teaches a third possible meaning for “iteratively traversing,” wherein multiple movements *and* repetition of the adjustment process would be required. In particular, the ’995 patent provides examples where “traversal” involves carrying out the adjustment sequence on a first node and then iterating the sequence on a second node (*i.e.*, at a next higher level in a hierarchy of resources) and iterating again on a third node (*i.e.*, at the same or a next higher level in the hierarchy of resources). Ex. B-2 at ¶ 28. Specifically, the ’995 patent teaches that “the resource allocation diagnosis algorithm *traverses the resource hierarchy*, starting from the bottom VM layer, *adjusting shares, limits and reservations of each layer.*” ’995 patent at 13:6–13 (emphasis added). The ’995 patent teaches that the adjustment process “can be also applied to the target VM’s parent RP node *and* other ancestor nodes higher up in the RP tree as necessary,” *id.* at 27–30 (emphasis added), including performing the adjustment process for “VM 420-3, its parent RP 406 and/or its next higher ancestor RP 404” or the parent’s “sibling RP 408,” *id.* at 15:33–41. These descriptions would indicate to a POSA that “iteratively traversing a resource hierarchy” might require both repeatedly moving along the hierarchy (*i.e.*, to multiple other nodes) and performing the adjustment procedures for each node. Ex. B-2 at ¶ 28.

A POSA would recognize the specification to provide multiple meanings for “iteratively traversing,” but would not be able to reasonably ascertain which applies. Ex. B-2 at ¶ 29. Because this term has three materially different meanings in view of the specification and “no informed and confident choice is available among the contending definitions,” the term is indefinite. *See HZNP Medicines LLC v. Actavis Lab ’ys UT, Inc.*, 940 F.3d 680, 698 (Fed. Cir.

2019) (citing *Media Rts. Techs.*, 800 F.3d at 1371 (citing *Nautilus*, 572 U.S. at 911 n.8)). Thus, “iteratively traversing a resource hierarchy” is not “precise enough to afford clear notice of what is claimed.” *Nautilus*, 572 U.S. at 909 (citing *Markman v. Westview Instruments, Inc.*, 517 U.S. 370, 373 (1996)).

“The fact that [VMware] can articulate a definition supported by the specification” by stringing together several dictionary definitions for its plain and ordinary meaning and alternate constructions does not save the term from indefiniteness. *Halliburton Energy Servs., Inc. v. M-I LLC*, 514 F.3d 1244, 1251 (Fed. Cir. 2008). Dictionary definitions “do not necessarily reflect the inventor’s goal of distinctly setting forth his invention,” *Phillips*, 415 F.3d at 1322, including here, where VMware’s dictionary definitions (and its analysis) do not account for other, different meanings communicated by the patent’s specification. *Anderson v. Int’l Eng’g & Mfg., Inc.*, 160 F.3d 1345, 1348–49 (Fed. Cir. 1998) (“A word describing patented technology takes its definition from the context in which it was used by the inventor.”).

VMware’s alternative construction should be rejected for the additional reason that it seeks to impermissibly broaden the independent claims through the principle of claim differentiation. “[C]laim differentiation is a rule of thumb that does not trump the clear import of the specification,” *Edwards Lifesciences LLC v. Cook Inc.*, 582 F.3d 1322, 1332 (Fed. Cir. 2009), and it “cannot broaden claims beyond their correct scope,” *Curtiss-Wright Flow Control Corp. v. Velan, Inc.*, 438 F.3d 1374 (Fed. Cir. 2006).

VMware argues that the independent claims should be construed to broadly encompass “moving down, up, or within a layer of the resource hierarchy” because certain dependent claims “limit ‘iteratively traversing a resource hierarchy’ by requiring analysis of an entity in the bottom layer of the hierarchy and then subsequent analysis of each parent resource of the entity.” *Supra*

§ III.A.2.a at pages 91-92. However, the specification does not support VMware’s overly broad construction of the independent claims. The patent does not describe or contemplate “moving down” a resource hierarchy, or “moving . . . within a layer of the resource hierarchy” without first traversing “higher.” *See* ’995 patent at 13:6–10 (“starting from the bottom layer”), 13:26–30 (“these procedures can be also applied to the target VM’s parent RP node and other ancestor nodes higher up in the RP tree.”), 13:36–37 (“Subsequently, for each parent of the current node . . .”), 15:28–44 (“traversing higher in the RP tree of a resource cluster and adjusting neighbor resource pools”), *id.* (“adjusting the . . . parameters of VM 420-3, its parent RP 406 35 and/or its next higher ancestor RP 404”), *id.* (“traverses higher up in the RP tree to the sibling RP 408 of the RP 406”). Nor does the specification indicate that “iteratively traversing a resource hierarchy” involves anything other than moving higher in the hierarchy or moving laterally after first moving higher. *See Retractable Techs., Inc. v. Becton, Dickinson & Co.*, 653 F.3d 1296, 1305 (Fed. Cir. 2011) (limiting construction of “body” “to what the specifications indicate the inventor actually invented” where “the specifications do not . . . indicate that the body is anything other than a one-piece body”). That some dependent claims require analyzing a resource at a “bottom layer” and subsequently analyzing “each parent” does not mean the independent claims necessarily encompass undisclosed downward traversal or lateral (“within a layer”) traversal without first traversing “higher.” *Id.* (“any presumption created by the doctrine of claim differentiation will be overcome by a contrary construction dictated by the written description or prosecution history”) (quotation omitted).

c. VMware’s Reply Position

Cirba’s indefiniteness argument focuses solely on the phrase “iteratively traversing.” It concedes that “resource hierarchy” has a plain and ordinary meaning in the art, *i.e.*, resources

organized in a hierarchy, such as in ranks, layers, or a tree structure. (*Supra* § III.A.2.b at page 92.)

Although Cirba asserts that “iteratively traversing” “is not [a] term of art in the relevant field” (Ex. B-2 ¶ 23), it also concedes that the plain meanings of “iteratively” (*i.e.*, repeatedly) and “traversing” (*i.e.*, “mov[ing] or pass[ing] along or through”) apply to the ’995 patent. (*Supra* § III.A.2.a at page 91-92.) Namely, Cirba agrees that the specification teaches that “‘iteratively traversing’ means simply ‘repeatedly moving’ along the hierarchy, *i.e.*, making more than one move among nodes (*e.g.*, node-to-node-to-node).” (*Compare supra* § III.A.2.b at page at 92-93 with *supra* § III.A.2.a at pages 91-92; *see also* Ex. A-20 ¶¶ 12-16.) This should end the inquiry.

VMware’s alternative construction does not seek to broaden the claims impermissibly, as Cirba claims. Nowhere does the patent limit “iteratively traversing” to “moving higher in the hierarchy or moving laterally after first moving higher,” as Cirba asserts. (*Supra* § III.A.2.b at page 96-97.) In an embodiment, the patent illustrates “iteratively traversing” by starting from the bottom layer in the resource hierarchy and moving up. (’995 patent, 17:49-53.) But that embodiment cannot limit “iteratively traversing” to any particular sequence of moves along the hierarchy. (“Although specific embodiments of the invention have been described and illustrated, the invention is not to be limited to the specific forms or arrangements of parts so described and illustrated.”).) *See Hill-Rom Servs., v. Stryker Corp.*, 755 F.3d 1367, 1371-72 (Fed. Cir. 2014) (disclosure of embodiment does not limit claim scope absent manifest exclusion or restriction).

Cirba also tries to manufacture an indefiniteness argument, alleging that the ’995 patent discloses two other “possible meanings”: (i) “the adjustment operations are performed at more than one node,” *i.e.*, “repeating the adjustment process”; and (ii) “repeatedly moving along the

hierarchy and performing the adjustment operation at each node.” (*Supra* § III.A.2.b at page 93-95.)

Cirba’s assertion that “iteratively traversing” could mean “repeating the adjustment process” (*id.* at 93-94) contradicts the plain language. Claims 1, 9, and 17 recite “generating a resource allocation recommendation based on the target resource allocation and the snapshot by *iteratively traversing* a resource hierarchy[.]” (Emphasis added.) The claims expressly require “iteratively traversing” a resource hierarchy—not “iteratively traversing *and adjusting*” or “iteratively *adjusting*” a resource hierarchy. A POSITA would not understand “traversing” to mean or include “adjusting” (Ex. A-20 ¶¶ 17-18), and Cirba cites no contrary evidence.

The specification does not help Cirba, as it never equates iteratively “traversing” and “adjusting.” (*Id.* ¶ 19.) In fact, the “resource allocation diagnosis algorithm” embodiment on which Cirba relies on confirms its plain meaning. (*Supra* § III.A.2.b at pages 93-94.) In the specification, the algorithm iteratively “traverses the resource hierarchy” by moving along the hierarchy more than once, “starting from the bottom VM layer.” (’995 patent, 13:6-10.) The algorithm moves from a client VM to the “VM’s parent RP node and other ancestor nodes higher up in the RP tree as necessary.” (*Id.*, 13:27-30.) In an example of the embodiment, the algorithm traverses “all the nodes” in a resource cluster. (*Id.*, 12:66-13:1, 13:39-41; Ex. A-20 ¶ 19.)

Although the specification discusses “three ordered procedures” (“AdjustLimit,” “AdjustShares,” and “AdjustReservations”), the specification does not mandate that iteratively “traversing” include this three-step “adjustment.” (Ex. A-20 ¶ 20.) Rather, the disclosed “adjustments” are simply exemplary ways of achieving a resource allocation recommendation based on a target resource allocation. (*Id.*) For example, “[i]n the ‘AdjustLimit’ procedure, the

limit parameter of a target VM is increased to be at least as large as the requested allocation.”
 (’995 patent, 13:13-21.) “In the ‘AdjustShares’ procedure, the requested resource allocation . . .
 is then used to determine the amount by which current shares have to be increased.” (*Id.*, 13:21-
 25.) And “[i]n the ‘AdjustReservations’ procedure, the reservation of the target VM is increased
 to be at least as large as the requested allocation.” (*Id.*, 13:25-27.)

Cirba’s argument that “iteratively traversing” also could mean “repeatedly moving along
 the hierarchy and performing the adjustment operation at each node” (*supra* § III.A.2.b at pages
 93, 95) fails for the same reason. Whether “adjustments” also could be performed when
 “iteratively traversing” is of no moment; the claims recite “iteratively traversing,” not
 “iteratively adjusting.” The claims do not require “performing [an]adjustment operation at more
 than one node,” and it would be improper to import that requirement into the claims. (Ex. A-20
 ¶ 21.)

d. Densify’s Sur-Reply Position

VMware fails to show how its proposed construction provides reasonable certainty as to
 which of the at least three possible ways a resource hierarchy may be traversed: (1) repeatedly
 “moving” along the hierarchy by moving among multiple nodes; (2) repeatedly “adjusting”
 operations at nodes on the hierarchy in addition to a single movement between two nodes; or (3)
 repeated movements *and* repeated adjustments. *Supra* § III.A.2.b at pages 92-95. The first
 possibility overlaps with VMware’s proposed construction. The second possibility is supported
 by disclosures that VMware ignores when it argues that “iteratively traversing” cannot also mean
 a single movement on the hierarchy with repeated “adjustments.” The specification describes a
 resource allocation diagnosis algorithm that “traverses the resource hierarchy” by starting from
 the bottom layer and “repeat[ing]” an “adjusting process . . . until the requested allocation has
 been achieved or new recommendations cannot be made.” ’995 patent at 13:6-30. Because this

can be accomplished in a *single* movement from one node to another node, as long as there are repeated (iterated) “adjustments” at the nodes, it presents a second distinct way the claim limitation can be met. In other words, the “traverse” requirement is satisfied by a single “movement” on the resource hierarchy, but the “iteratively” requirement is met by repeated “adjustments.” The third possibility is simply a combination of the first two possibilities. The claim language does not give a POSA reasonable certainty as to what the term requires. Ex. B-2 ¶¶ 23-29.

3. “target resource allocation” (’995 Patent: claims 1, 9, and 17)

VMware’s Proposed Construction	Densify’s Proposed Construction
This term should not be construed as it takes its plain and ordinary meaning.	“desired resource allocation”

a. VMware’s Opening Position

This phrase should not be construed, as it takes its plain and ordinary meaning. “Target” is a common English word meaning a goal or an objective. (*See* Ex. A-7 (*New Oxford Am. Dictionary* (2d ed. 2005)) at 1727 (defining “target” as “an objective or result toward which efforts are directed”); Ex. A-6 (*Merriam-Webster’s Collegiate Dictionary* (11th ed. 2007)) at 1278 (defining “target” as “a goal to be achieved”).) The specification supports giving “target” its ordinary English meaning without construction. In one disclosed embodiment, the specification expressly describes an exemplary target resource allocation as a “resource allocation goal,” specifically “the current resource demand of a VM.” (’995 patent, 11:43-48.)

Cirba merely wishes to replace one common English word, “target,” with a different word, “desired.” And as Cirba’s proposed construction concedes, the phrase “resource allocation” requires no construction.

b. Densify’s Answering Position

When the term “target” is used to modify “resource allocation,” the patentees act as their own lexicographer, using it synonymously with “desired.” ’995 patent at 1:14–17 (“Resource allocation techniques. . . are important to ensure that [] clients are operating at *desired or target* levels.”) (emphasis added).

As VMware’s own extrinsic evidence shows, “target” has a variety of plain and ordinary meanings, including “a mark or point at which someone fires or aims” (*n.*) or “aim or direct” (*v.*). To guard against unnecessary confusion, and to ensure the patentees’ lexicography is properly applied, “target” should be as construed as “desired,” per the specification’s use of the term.

c. VMware’s Reply Position

The claims’ use of “target” is consistent with its ordinary meaning as a goal. The patent describes an exemplary target resource allocation as a “resource allocation goal,” such as “the current resource demand of a VM.” (’995 patent, 11:43-48.)

Relying on the patent’s Background, Cirba tries to improperly change the claim scope by substituting “desired” for “target.” (*Supra* § III.A.3.b.) But the Background does not use “target” and “desired” synonymously. (’995 patent, 1:16-17 (“clients are operating at *desired or target* levels”) (emphasis added).) The claims also use the word “target,” not “desired,” which connotes a user’s subjective wishes. Cirba’s construction could confusingly suggest that the jury must assess the subjective intentions of some individual to evaluate infringement. The Court should reject Cirba’s attempt to swap out the ordinary word “target” for the subjective “desired.”

d. Densify’s Sur-Reply Position

VMware states that the patent describes an exemplary target resource allocation as a “resource allocation goal” meaning it is a *desired* resource allocation that an individual wants to

achieve. *Supra* § III.A.3.c. The specification supports that “target” and “desired” are synonymous when the individual wants to achieve this “goal.” ’995 patent, 11:43–48.

Moreover, a jury would understand the differences between target, desired, and a goal since the patentees’ lexicography applied equally to “target” and “desired” per the specification’s use of the term. *Id.* at 1:14–17. These are common terms that most people use in their everyday vernacular. There is no confusion in Densify’s proposed construction.

4. “resource configuration action” (’995 Patent: claims 1, 9, and 17)

VMware’s Proposed Construction	Densify’s Proposed Construction
This term should not be construed as it takes its plain and ordinary meaning.	“action relating to resource configuration in accordance with resource allocation recommendation”

a. VMware’s Opening Position

This phrase also should not be construed, as it takes its plain and ordinary meaning. Cirba effectively concedes this by recasting the phrase “resource configuration action” as an “action relating to resource configuration.” But Cirba then improperly adds the language “in accordance with [the] resource allocation recommendation.” Cirba’s added language is improperly limiting, as the claims do not require that the “resource configuration action” be “in accordance with [the] resource allocation recommendation.”

b. Densify’s Answering Position

The term “resource configuration action” does not have an accepted meaning in the art. The ’995 patent uses the term in the context of “resource allocation recommendations” which are “based on a target resource allocation and the snapshot.” ’995 patent at 8:65-67. The recommendation “may specific at least one *resource configuration action* or at least one capacity expansion action ... to meet the target resource allocation” (*id.* at 8:67-9:4) (emphasis added)

and then lists the possible actions the recommendation may recommend: “maintain the current configurations,” “move one or more clients from their current host computers . . . to other host computers,” “power down one or more clients or host computers,” “change the resource entitlement for one or more host computer.” *Id.* 9:4-18. A POSA would have understood the disclosure to tie the scope of the actions to the recommendation results, because that is the only context in which the term is used. In other words, the claimed actions must be “in accordance with the resource allocation recommendation.” That the literal text of the claims does not include those words is not controlling. *See Irdeto Access, Inc. v. Echostar Satellite Corp.*, 383 F.3d 1295, 1300 (Fed. Cir. 2004) (“[A]bsent such an accepted meaning, we construe a claim term only as broadly as provided for by the patent itself.”).

c. VMware’s Reply Position

The plain claim language requires only that the resource allocation recommendation “*specif[y]* at least one resource configuration action or at least one capacity expansion action[.]” (’995 patent, claims 1, 9, and 17 (emphasis added).) *I.e.*, if the resource allocation recommendation “specifies” the resource configuration action, the limitation is met.

The plain claim language does not state that the resource configuration action is “in accordance with” the resource allocation recommendation, as Cirba proposes. Cirba does not explain why it is necessary to insert “in accordance with” into the claims, which already state that the resource allocation recommendation “specifies” the resource configuration action. Cirba’s proposed construction also suffers from two deficiencies. First, “in accordance with” is vaguer than the straightforward “[is] specifi[ed].” Second, it is potentially narrower, as it implies that performing the specified “resource configuration action” depends on other unidentified requirements in the “recommendation.”

Cirba’s specification excerpts do not support its construction. They identify exemplary actions that the resource allocation recommendation can specify, but do not require that the claimed “resource configuration action” be “in accordance with” the resource allocation recommendation. (*See supra* § III.A.4.b.) The same is true of dependent claims 3, 5, and 6, which also recite exemplary “resource configuration actions”—but also do not require that they be “in accordance with” the recommendation.

d. Densify’s Sur-Reply Position

“Resource configuration action” has no accepted meaning in the art, and VMware offers no serious argument to the contrary. *Supra* § III.A.4.b. Densify’s proposed construction explains how the required action must – as taught by the specification – be in accordance with the “resource allocation recommendation.” ’995 patent at 8:65-67; 8:67-9:4; and 9:4-18. No other kind of resource configuration action is disclosed or suggested in the intrinsic record.

5. “capacity expansion action” (’995 Patent: claims 1, 9, and 17)

VMware’s Proposed Construction	Densify’s Proposed Construction
This term should not be construed as it takes its plain and ordinary meaning.	“action relating to capacity expansion in accordance with resource allocation recommendation”

a. VMware’s Opening Position

This phrase also should not be construed, as it takes its plain and ordinary meaning. Here, too, Cirba effectively concedes as much by recasting “capacity expansion action” as “action relating to capacity expansion.” And here, too, Cirba’s added language is improperly limiting, as the claims do not require that the “capacity expansion action” be “in accordance with [the] resource allocation recommendation.”

b. Densify's Answering Position

As with “resource configuration action,” the term “capacity expansion action” does not have an accepted meaning in the art. It too is defined by the '995 patent as a set of actions relating to capacity expansion that may be included in a “recommendation” resulting from a “resource allocation diagnosis.” See '995 patent at 8:65–9:19. The recommendation “may specific at least one resource configuration action or at least one *capacity expansion action* ... to meet the target resource allocation.” *Id.* 8:67-9:4 (emphasis added). For instance, the patent indicates that “[i]f the current resource entitlement does not match the required resource allocations, the algorithm increases or recommends increasing the capacity of the resource cluster, at step 624.” *Id.* at 15:17–20. The patent also indicates that when “the resource allocation diagnosis algorithm cannot find a set of recommendations” due to a small cluster size, it “suggests increasing the cluster capacity.” *Id.* at 15:55–62. A POSA would have understood the disclosure to tie the scope of the actions to the recommendation results, such that the claimed actions must be “in accordance with the resource allocation recommendation.”

c. VMware's Reply Position

Here, too, Cirba improperly limits the claims by requiring that the capacity expansion action be “in accordance with” the resource allocation recommendation. As with “resource configuration action,” nothing in the claims or specification requires this.

d. Densify's Sur-Reply Position

Like “resource configuration action,” VMware's argument does not help the jury understand the term “capacity configuration action,” which likewise has no accepted meaning in the art. “In accordance with resource allocation recommendation” provides the only disclosed context for what the “action” must be. '995 patent at 8:65-67; 8:67-9:4; 15:17-20; and 15:55-62.

B. U.S. Patent No. 9,766,945 (“’945 Patent)**1. “calculating a target resource configuration” (’945 Patent: claims 1, 8, and 15)**

VMware’s Proposed Construction	Densify’s Proposed Construction
This term should not be construed as it takes its plain and ordinary meaning.	“calculating a desired resource configuration”

a. VMware’s Opening Position

The ’945 patent (D.I. 1032-4, Ex. L) describes methods for scheduling computing resources in datacenters, so that datacenter “[r]esources can be optimized across layers” of hosts, virtual machines, and containers. (*Id.*, 2:8-13.) Hosts are configured to provide a virtualization layer of computing resources shared by one or more virtual machines. (*Id.*, 2:22-31.) In turn, containers may be executed in VMs running on the hosts. (*Id.*, 2:57-64.)

Containers are “software instances that enable virtualization at the operating system level.” (*Id.*, 2:64-66.) They “appear as unique servers from the standpoint of an end user that communicates with the containers.” (*Id.*, 3:1-3.) “However, from the standpoint of the operating system that manages the host computer on which the containers execute, the containers are user processes that are scheduled and dispatched by the operating system.” (*Id.*, 3:4-7.) Although a container can execute either in a host or a virtual machine, at least one container is “allocat[ed] . . . to be executed in” a VM in the claimed inventions (*see id.*, 12:10-12).

The phrase “‘calculating a target resource configuration’ for one or more VMs” in claims 1, 8, and 15 needs no construction, as it takes its plain meaning in the art and in view of the intrinsic evidence. It is not indefinite.

Claims 1, 8, and 15 use the words in this phrase in their ordinary sense. The words “calculating” and “target” are common. “Calculating” means “ascertain[ing] by computation;

reckon[ing]” (Ex. A-8 (*Am. Heritage Dictionary of the English Language* (5th ed. 2011)) at 263), or “determin[ing] or ascertain[ing] by mathematical methods; comput[ing]” (Ex. A-9 (*Random House Unabridged Dictionary* (2d ed. 1993)) at 295). A “target” is “an objective or result toward which efforts are directed” (Ex. A-7 (*New Oxford Am. Dictionary* (2d ed. 2005)) at 1727), or “a goal to be achieved” (Ex. A-6 (*Merriam-Webster’s Collegiate Dictionary* (11th ed. 2007)) at 1278). In the art, a “configuration” is “the way in which a system, whether it be hardware and/or software, is set up” (Ex. A-10 (*Collins Dictionary of Science* (2005)) at 143), and a “resource” is “[a]ny part of a computer system or a network, such as a disk drive, printer, or memory, that can be allotted to a program or a process while it is running” (Ex. A-1 (*Microsoft Computer Dictionary* (5th ed. 2002)) at 451).²¹ Putting these definitions together, the phrase “calculating a target resource configuration” means determining by computation a goal for the set-up [for VMs]. But the original language of the phrase, without construction, conveys this same meaning clearly.

The phrase’s use in the claims also readily apprises one skilled in the art of its meaning. In claim 1, for example, the phrase appears as a step in a multi-step process. The first step — “determining . . . a resource allocation for one or more virtual machines (VMs), and a resource usage for one or more containers” — provides input data for subsequent steps. The third step (at issue here) is “‘calculating a target resource configuration’ for one or more VMs.” The fourth and sixth steps then recite what is done with the calculated target resource configuration, viz., “removing or adding resources to each of the one or more VMs *for which a target resource configuration was calculated to achieve the target resource configuration for each VM,*” and

²¹ A POSITA understands that the term “resource” includes CPU, memory, or storage. (*See, e.g.,* ’945 patent, 2:22-50, 4:22-39, 5:50-60.)

“allocating the one or more containers to be executed in the one or more VMs *based on the resource configuration of each VM* and the resource usage of the containers.” (Emphasis added.)

So, too, the specification’s discussion of “calculating a target resource configuration” provides ready context for its plain meaning. The specification explains that, “[a]s containers are started and shut down on various VMs, some VMs may end up with more resources than necessary for their assigned containers, while other VMs become over-committed, with not enough resources.” (’945 patent, 1:29-34.) Accordingly, “embodiments described herein optimize hardware resources by providing a correct resource allocation to host VMs by looking at the consumption of containers.” (*Id.*, 4:7-10.) In a specific example, a target resource configuration for a VM of 3 GB is calculated by equating it to the container resource usage: the containers need 3 GB of resources, and the VM is then “right-sized by . . . removing resources” to achieve the calculated configuration by reducing the VM’s memory from 5 GB to 3 GB. (*Id.*, 5:26-37, FIG. 5, claim 1 (the third and fourth steps).)

b. Densify’s Answering Position²²

The term “calculating a target resource configuration” does not have an accepted meaning in the art. The only description in the specification of the ’945 patent of any calculation of a resource configuration is directed to the calculation of an “ideal” configuration. ’945 patent at 6:43-58. As VMware’s own extrinsic evidence shows, “target” has a variety of plain and ordinary meanings, including “a mark or point at which someone fires or aims” (*n.*) or “aim or direct” (*v.*). To guard against unnecessary confusion, “target” should be as construed as “desired.”

²² Densify withdraws its initial construction of “indefiniteness” for this term and proposes the construction herein.

c. VMware's Reply Position

Without intrinsic support or expert opinion, Cirba again proposes to swap out the ordinary word “target” for the subjective word “desired,” which does not appear in the ’945 patent. Cirba’s construction ignores the plain language and the specification’s discussion and example of calculating a target resource configuration. (*Supra* § III.B.1.a at pages 109-111; ’945 patent, 5:26-37, FIG. 5, claim 1.) And as discussed for the ’995 patent above, “desired” confusingly connotes a subjective wish of the user. Cirba’s attempt to change the claims’ meaning should be rejected.

d. Densify's Sur-Reply Position

The term “target” has multiple possible meanings. Densify’s proposed construction tracks the specification and eliminates confusion.

C. U.S. Patent No. 10,025,638 (“’638 Patent”)

1. “operationally distinct cloud computing facilities” (’638 Patent: claims 10 and 19)

VMware's Proposed Construction	Densify's Proposed Construction
Not indefinite. This term should not be construed as it takes its plain and ordinary meaning.	Indefinite for failure to point out with particularity and distinctly claim the subject matter such that one of ordinary skill in the art would be reasonably apprised of the bounds of the asserted claims.

a. VMware's Opening Position

The inventions disclosed in the ’638 patent (D.I. 1032-4, Ex. M) improve the management of cloud-computing facilities that include virtual data centers. The patent describes methods that provide services across multiple cloud-computing facilities, such as transfer of a VM from one cloud-computing facility to another. (*Id.*, FIG. 10, 12:33-13:2, 17:65-18:5, 20:1-

37.) The disclosed inventions allow a system administrator to manage a cloud-aggregation that includes multiple remote cloud-computing facilities. (*Id.*, 12:61-65, 18:45-19:12.)

The asserted independent claims recite “operationally distinct cloud computing facilities” as structural elements that are used, respectively, in the steps of a method (claim 10) and in computer instructions that encode a method (claim 19). The phrase requires no construction, as its words take their plain and ordinary meaning, and it is not indefinite.

The word “operationally” means functionally, and the word “distinct” means separate. (*See* Ex. A-6 (*Merriam-Webster’s Collegiate Dictionary* (11th ed. 2007)) at 869 (defining “operate” as “perform a function,” “operation” as “the quality or state of being functional or operative,” and “operational” as “of or relating to operation or to an operation . . . — operationally *adv*”); Ex. A-11 (*Collins English Dictionary* (11 ed. 2011)) at 1161 (defining “operate” as “function or cause to function,” “operational” as “of or relating to an operation or operations . . . > operationally *adv*,” and “distinct” as “separate”).) A POSITA therefore would understand “operationally distinct cloud computing facilities” to refer to cloud computing facilities that function separately from one another. But that express construction is unnecessary, as the phrase uses ordinary words with ordinary meanings.

The ’638 patent specification confirms this plain meaning of “operationally distinct cloud computing facilities” and that one skilled in the art would readily understand it. As the specification explains, the inventions seek to provide “useful and powerful distributed operations between remote virtual data centers within different remote physical cloud-computing facilities.” (*Id.*, 20:1-22.) The inventions also enable “an administrator [to] manage an entire cloud aggregation comprising many different remote virtual data centers running within many different remote cloud-computing facilities.” (*Id.* at 18:45-51; *see also id.*, 18:14-31 (describing

exemplary “centrally managed” cloud-computing facilities); 20:23-37 (“a cloud aggregation [] can be centrally managed . . . and [] can employ cloud-computing facilities distributed across the face of the earth”).) The cloud-computing facilities that are subject to centralized management function separately from one another. “In general, the cloud-computing facilities that together form a multiple-cloud-computing aggregation through distributed services provided by the VCC server and VCC nodes are geographically and operationally distinct.” (*Id.*, 12:65-13:2.)

Excerpted and annotated FIG. 10 depicts seven operationally distinct cloud-computing facilities (1002 to 1008) that function separately from one another, but are centrally managed:




FIG. 10

(*Id.*, 12:38-51.) Exemplary cloud-computing facility 1002 is a private multi-tenant cloud. (*Id.*, 12:39-43.) The remaining cloud-computing facilities, 1003-1008, are all shown as distinct facilities too. (*Id.*, FIG. 10, 12:38-39, 12:43-51.) Each is either a public or private cloud-computing facility that includes single-tenant virtual data centers (1003 and 1006) or multi-tenant virtual data centers (1004 and 1007-1008) or is a third-party cloud-services facility (1005). (*Id.*) Although FIG. 10 depicts different types of operationally distinct cloud-computing

facilities, the inventions also can include operationally distinct cloud-computing facilities of the same type. (*Id.*, 18:26-29 (“The multi-cloud aggregation may include only remote clouds of one type, such as clouds managed by VDC management servers, and thus may be homogeneous”).)


The prosecution history confirms that “operationally distinct cloud-computing facilities” takes its ordinary meaning, need not be construed, and is not indefinite. During prosecution, the patentee explained that:

[The cloud-computing facilities] are operationally distinct because . . . the virtual data centers [within the cloud-computing facilities] are themselves large distributed systems that include interfaces and abstraction layers that provide access to well-defined sets of services that are carried out entirely within the virtual datacenters. Typical virtual data centers do not provide services that are distributed across multiple virtual data centers.

(D.I. 1032-12, Ex. T at VMWCC00016944 (2016/5/4 Reply Brief at 5).) The patentee then stated the disclosed inventions involve an “aggregation of multiple cloud-computing facilities [that] provide services distributed across multiple cloud-computing facilities.” (*Id.* at VMWCC00016947 (2016/5/4 Reply Brief at 8).)

b. Densify’s Answering Position

The “operationally distinct” language of this claim element is not a term of art and the intrinsic record offers no clear indication of what it is intended to mean. The specification uses the term “operationally distinct” only once and in connection with Figure 10, stating that “the cloud-computing facilities that together form a multiple-cloud-computing aggregation through distributed services provided by the VCC server and VCC nodes are geographically and operationally distinct.” ’638 patent at 12:56-13:2; Ex. B-2 at ¶¶ 30-31.



The specification identifies 1002-1008 in Figure 10 each as “cloud-computing facilities,” explaining that “[c]loud computing facility 1002 is a private multi-tenant cloud with a cloud director 1010 that interfaces to a VDC management server 1012 to provide a multi-tenant private cloud comprising multiple tenant-associated virtual data centers.” ’638 patent at 12:38-43. While “cloud computing facility 1002” is expressly defined, the other cloud computing facilities have a variety of potential characteristics: they may be public, private, single-tenant, multi-tenant or “any of the various different kinds of third-party cloud-services facilities.” *Id.* at 12:43-51. None of these options sheds light on what it means for such facilities to be “operationally distinct.” Ex. B-2 at ¶ 32.

The prosecution history further confuses the situation. After receiving a Final Rejection, VMware filed an appeal to the Board. In the Examiner’s Answer to VMware’s Appeal Brief, the Examiner explained his own confusion regarding the term “operationally distinct”:

As specified in the claims, the distinction necessary between the cloud computing facilities is that they are “*operationally distinct*”,

and while *no specific delineation is set forth in the claims*, the examiner relies on the description in paragraph [0044] of the applicant's specification which is itself vague at best. *As best the examiner can determine, it is only necessary that the node operate independent of each other in some way (and not even necessarily be geographically distinct as the applicant provides that a separate geographical distinction)*. With that in mind, the examiner affirms the teaching of Martin anticipates operational distinct nodes in that each virtual hypervisor layer described in paragraph [0036] is in and of itself operationally distinct from the other hypervisor layers, which is the nature of virtual machines.

Ex. B-3 ('638 File History, March 4, 2016 Examiner's Answer) at 10 (emphasis added). Instead of confirming or correcting the Examiner's understanding of what "operationally distinct" meant, Applicant ignored it, and responded with a statement that appears to offer an alternate test for what qualifies as operationally distinct:

The claimed multiple cloud-computing-facility aggregation includes multiple, operationally distinct cloud-computing facilities. These are shown in Figure 10 as cloud-computing facilities 1002-1008. They are *operationally distinct because*, as explained in copious background material included in the current application, *the virtual data centers are themselves large distributed systems that include interfaces and abstraction layers that provide access to well-defined sets of services that are carried out entirely within the virtual data centers*.

D.I. 1032-13, Ex. T at 24 ('638 File History, May 4, 2016 Reply Br. at 5) (emphasis added).

Based on the Applicant's remarks during prosecution, a POSA would be left to ponder several questions. For example, are the claimed facilities required to include "abstraction layers" to be "operationally distinct?" Do those abstraction layers have to "provide access to well-defined sets of services?" Do the services have to be "carried out entirely within the virtual data centers?" And importantly, is there a requirement that they be "geographically distinct,"— as the Examiner theorized (but expressed doubt about)? In ultimately resolving the appeal in Applicant's favor, the Board *did not* decide the issue of what qualifies as "operationally distinct" in the context of

the '638 patent, or provide any further comment on the Examiner's or Applicant's respective positions concerning the claim. Ex. B-2 at ¶ 33.

A POSA therefore would have been left with no clear understanding of what is required for cloud computing facilities to be “operationally distinct” because the intrinsic record provides no reasonable certainty as to what that means, including whether such facilities must be “geographically distinct” from one another or have the particular qualities articulated in Applicant's argument. Given the confused file history, a POSA would not understand what it means for a computing facility to be “operationally distinct.” *HZNP Medicines LLC v. Actavis Lab'ys UT, Inc.*, 940 F.3d 680, 698 (Fed. Cir. 2019) (“[A] claim is indefinite if its language might mean several different things and no informed and confident choice is available among the contending definitions.” (quoting *Media Rights Techs., Inc. v. Capital One Fin. Corp.*, 800 F.3d 1366, 1371 (Fed. Cir. 2015))); *Nautilus*, 572 U.S. at 901 (“[A] patent is invalid for indefiniteness if its claims, read in light of the *specification delineating the patent, and the prosecution history*, fail to inform, with reasonable certainty, those skilled in the art about the scope of the invention.” (emphasis added)); Ex. B-2 at ¶ 34.

VMware suggests that the Court should rely on non-technical dictionary definitions of the words “operation” and “distinct” and synthesize its own definition or, more problematically, leave the jury to engage in that exercise on its own. *Supra* § III.C.1.a at page 111. But even assuming, as VMware suggests, “distinct” means “separate,” what does it mean to be “operationally” separate? The Examiner called out the “vague” nature of the specification's disclosures regarding this term and provided the patentee an opportunity to affirm or correct what he surmised the term to mean. But patentee did not do so, leaving scope of the claim not

reasonably ascertainable to a POSA. Ex. B-2 at ¶ 35. Accordingly, the term “operationally distinct cloud computing facilities” is indefinite.

c. VMware’s Reply Position

Conceding that “cloud computing facilities” is definite, Cirba focuses its indefiniteness argument solely on the modifier “operationally distinct.” Nothing about the ordinary words “operationally” and “distinct” or their combination with the phrase “cloud computing facilities” is indefinite. Cirba does not dispute that “operationally” means functionally, and “distinct” means separate. (*Compare supra* § III.C.1.a at page 111 *with supra* § III.C.1.b at page 115-16.) A POSITA understands that “operationally distinct” cloud computing facilities simply refers to cloud computing facilities “that function separately from one another.” (Ex. A-19 ¶¶ 11-13.)

The specification supports this plain meaning. (*Id.* ¶¶ 14-15.) It describes providing centralized management to deliver distributed services across multiple cloud-computing facilities that otherwise function separately from one another. (*Id.*; *supra* § III.C.1.a at pages 111-113.) As a result, the claimed inventions can provide “useful and powerful distributed operations,” including for example, (i) a cloud computing facility “to be backed up to” another distinct cloud computing facility; (ii) “high-bandwidth secure-communications links” between two cloud computing facilities; and (iii) “a full or partial failover” of a cloud computing facility to another to extend high-availability computing. (’638 patent, 20:1-22; Ex. A-19 ¶ 14.) Cirba still ignores Figure 10’s clear illustration of cloud computing facilities that function separately from one another and the specification’s accompanying discussion, which VMware previously noted. (*Supra* § III.C.1.a at page 111-113; Ex. A-19 ¶ 15.)

Although Cirba argues that “[t]he specification uses the phrase ‘operationally distinct’ only once” (*supra* § III.C.1.b at page 113), “[t]here is no requirement that every term appearing in the claims must be specifically defined or used in the specification.” *Magna Elecs., Inc. v.*

TRW Auto. Holdings Corp., No. 12-cv-654, 2015 WL 11401855, at *14 (W.D. Mich. Apr. 28, 2015); *see also In re Skvorecz*, 580 F.3d 1262, 1268-69 (Fed. Cir. 2009).

Concocting an indefiniteness argument, Cirba asserts that the POSITA would not understand whether “operationally distinct” also means “geographically distinct.” (Cirba Resp. Br. at 14). Although the specification discloses that the cloud computing facilities can be “geographically distinct” and “operationally distinct,” Cirba ignores that the patent describes the two concepts separately. (’995 patent, claims 10, 19; 12:65-13:2 (“the cloud-computing facilities that together form a multiple-cloud-computing aggregation . . . are geographically and operationally distinct.”); Ex. A-19 ¶ 16.) The claims require only the latter.

Cirba’s reliance on the prosecution history also is misplaced. This phrase “operationally distinct cloud computing facilities” is presumed to be definite, *see S3 Inc. v. NVIDIA Corp.*, 259 F.3d 1364, 1367 (Fed. Cir. 2001), and the PTO did not determine otherwise. Although the Examiner sought clarification about the phrase “operationally distinct,” the applicant provided it, explaining that “operationally distinct cloud computing facilities” refers to cloud computing facilities that provide access to services that are carried out entirely within the facilities. (D.I. 1032-13, Ex. T at VMWCC00016944 (2016/5/4 Reply Br. at 5).) The applicant’s explanation is consistent with the plain meaning of “operationally distinct cloud computing facilities” as facilities that function separately from one another. (Ex. A-19 ¶¶ 17, 18.)

Moreover, the PTO Board’s decision—finding that “the Examiner has not shown how [the cited prior art] discloses the claimed ‘multiple, operationally distinct cloud-computing facilities’”—confirms that “operationally distinct” is not indefinite. (D.I. 1032-13, Ex. T at VMWCC00016966 (2017/11/8 Board Decision at 5); Ex. A-19 ¶19.) The Board plainly “had no trouble analyzing the term in relation to [the] prior-art.” *3G Licensing, S.A. v. Blackberry Ltd.*,

No. 17-82-LPS-CJB, 2018 WL 4375091, at *8 (D. Del. Sept. 13, 2018); *see also Sonix Tech. Co. v. Publ'ns Int'l, Ltd.*, 844 F.3d 1370, 1379-80 (Fed. Cir. 2017) (claim not indefinite where PTO had no difficulty applying limitation to prior art).

d. Densify's Sur-Reply Position

VMware defines “operationally distinct cloud computing facilities” to mean “cloud computing facilities ‘that function separately from one another’” (*Supra* § III.C.1.c at page 117) **and also** “cloud computing facilities that provide access to services that are carried out entirely within the [cloud computing] facilities” (*id.* at 118) (citing applicant’s explanation of the phrase to the PTO during prosecution). VMware’s own alternative definitions—which are not co-extensive and have different scopes—prove the indefinite nature of the claim term.

VMware’s first definition (“cloud computing facilities ‘that function separately from one another’”) is itself indefinite. VMware provides no explanation of what it means for each cloud computing facility to “function separately.” The only specification disclosure it cites to suggests the opposite, teaching that a “multi-cloud aggregation” includes components that “**cooperate** to provide services that are distributed across multiple clouds.” ’638 patent, 12:33-38 (emphasis added). “Cooperation” connotes interoperation, not separately functioning. A POSA would wonder whether cloud facilities that “cooperate” and “provide services ... distributed across multiple clouds” are “functionally separate.” Ex. B-14 ¶¶ 19-22.

VMware’s second definition (“cloud computing facilities that provide access to services that are carried out entirely within the [cloud computing] facilities”) is at odds with its first definition. Whereas the specification teaches – and VMware specifically cites to – disclosures that characterize a multi-cloud aggregation (i.e., the recited “multiple, operationally distinct cloud computing facilities” of asserted claims 10 and 19) as “provid[ing] services ... **distributed across multiple clouds**,” VMware simultaneously contends that “operationally distinct” means

“cloud computing facilities that provide access to services that are carried out *entirely within the [cloud computing] facilities*.” These definitions are inconsistent, and a POSA would be left to wonder which controls and if (and how) they may be reconciled. Ex. B-14 ¶¶ 19-22.

VMware’s expert, Dr. Menascé, offers only conclusory opinions. *See* Ex. A-19 ¶¶ 17-20. His report quotes verbatim from the specification, but provides no analysis and attempts no reconciliation among what he purports to be the plain meaning of the term (“function separately”), what the specification teaches (provid[ing] services ... distributed across multiple clouds”), and what the applicant said the term means (“provide access to services that are carried out entirely within the [cloud computing] facilities”). *Id.* These non-co-extensive, inconsistent meanings would deprive a POSA of reasonable certainty as to what “operationally distinct cloud computing facilities” requires.

2. “cloud-connector server” (’638 Patent: claims 10 and 19)

VMware’s Proposed Construction	Densify’s Proposed Construction
<p>This term should not be construed as it takes its plain and ordinary meaning.</p> <p>Alternatively, this term should be construed as “a dedicated server computer or virtual machine running within a server computer.”</p>	<p>“a dedicated server computer or virtual machine running within a server computer that interfaces to users on a remote computer and that interfaces to remote cloud-connector nodes”</p>

a. VMware’s Opening Position

“Cloud-connector server” requires no construction. A “server” is well-known in the art as “a computer or program that responds to commands from a client.” (Ex. A-1 (*Microsoft Computer Dictionary* (5th ed. 2002)) at 474; *see also* Ex. A-4 (*Dictionary of Computing* (6th ed. 2010)) at 301 (defining “server” as a “dedicated computer which provides a function to a network”).) And “cloud-connector” is a descriptive label that (i) specifies the server that manages the multiple cloud-computing facilities and (ii) differentiates the “cloud-connector

server” from another “server”, a “virtual data center management server.” As the plain language of the independent claims further explains, the cloud-connector server “provides an electronic cloud-aggregation management interface.” (’638 patent, claims 10, 19.)

Because “cloud-connector server” has a plain and ordinary meaning to one skilled in the art that is apparent from the claims, the phrase is not indefinite. But the plain and ordinary meaning of the phrase also is apparent from the specification. The specification describes the “cloud-connector server” as (1) one or more physical servers or (2) a virtual machine or virtual machine appliance running within a virtual machine.²³ For example, the specification discloses that the “cloud-connector server” can be “one or more physical servers located within a private cloud-computing facility.” (*Id.*, 13:5-8.) Alternatively, the “cloud-connector server” may be “a virtual appliance within a VDC management server.” (*Id.*, 12:51-57.) “A virtual appliance is a software service that is delivered as a complete software stack installed within one or more virtual machines[.]” (*Id.*, 9:21-23.)

The prosecution history confirms the phrase’s plain, ordinary, and definite meaning. The patentee explained that “[t]he cloud-connector server is a server computer that runs a cloud-connector-server application within a cloud-connector virtual machine” . . . or “[a]lternatively, the cloud-connector server may run as a virtual machine in another type of server computer[.]” (D.I. 1032-12, Ex. T at VMWCC00016851 (2015/09/23 Appeal Brief at 10).) “The cloud-connector server is thus a dedicated server computer or virtual machine running within a server computer that interfaces to users on a remote computer and that interfaces to remote cloud-connector nodes.” (*Id.* at VMWCC00016852 (2015/09/23 Appeal Brief at 11).)

²³ The specification also references an exemplary “cloud-connector server” as the “virtual-cloud-connector server” or the “VCC server.”

b. Densify's Answering Position²⁴

“Cloud-connector server” is not a term known in the art. During prosecution, the patentee told the Patent Office that the term was a new phrase developed for the instant application:

Please note that the phrases “*cloud-connector server*” and “*cloud-connector node*” ***do not have well-known or accepted meanings within the computing arts***. These phrases were developed for the current application to describe the new cloud-computing-facility-aggregation components to which the current application, in part, is directed. These phrases must be interpreted according to the definitions and explanations provided in the current application.

Ex. B-4 ('638 File History, September 23, 2015 Appeal Brief) at 12. The patentee provided an express definition:

The *cloud-connector server* is thus ***a dedicated server computer or virtual machine running within a server computer that interfaces to users on a remote computer and that interfaces to remote cloud-connector nodes***.

Id. at 11. A POSA would rely on this clear definition of “cloud-computing server.” *See, e.g., CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1366 (Fed. Cir. 2002) (“[T]he claim term will not receive its ordinary meaning if the patentee acted as his own lexicographer and clearly set forth a definition of the disputed claim term in either the specification or prosecution history.”); *Honeywell Inc. v. Victor Co. of Japan*, 298 F.3d 1317, 1323 (Fed. Cir. 2002) (“It is well settled that a patentee may define a claim term either in the written description of the patent or, as in the present case, in the prosecution history.”); *Jack Guttman, Inc. v. Kopykake Enterprises, Inc.*, 302 F.3d 1352, 1361 (Fed. Cir. 2002) (“Because the patentee provided an explicit definition of “photocopy machine” in the specification, and because no other intrinsic

²⁴ Densify withdraws its initial construction of “indefiniteness” for this term and proposes the construction herein.

evidence casts doubt on that definition, the patentee’s definition controls.”); *Shire Dev. LLC v. Teva Pharms. USA, Inc.*, No. 1:17-CV-01696-RGA, 2019 WL 969638, at *10 (D. Del. Feb. 28, 2019) (construing term in accordance with patentee’s “lexicographical definition” from the specification).

c. VMware’s Reply Position

Cirba disregards VMware’s opening brief, which established that (1) a “server” is well-known in the art as “a computer or program that responds to commands from a client”; and (2) “cloud-connector” is a descriptive label for a particular server in the claims. (*Supra* § III.C.2.a at page 120-121.) As “cloud-connector server” has a plain and ordinary meaning to a POSITA that is apparent from the claims and specification, it does not need to be construed. In the alternative, the Court should adopt VMware’s construction of “a dedicated server computer or virtual machine running within a server computer.”

Cirba’s proposed additional language—“that interfaces to users on a remote computer and that interfaces to remote cloud-connector nodes”—ignores the patent’s broad description of the “cloud-connector server.” The specification describes it as (1) one or more physical servers or (2) a virtual machine or virtual machine appliance running within a virtual machine. (’638 patent, 13:5-8 (“cloud-connector server” can be “one or more physical servers located within a private cloud-computing facility”); 12:51-57 (“cloud-connector server” may be “a virtual appliance within a VDC management server”).) Cirba also ignores the applicant’s broad description of this server during prosecution. There, the applicant explained that “[t]he cloud-connector server is a server computer that runs a cloud-connector-server application within a cloud-connector virtual machine” or “may run as a virtual machine in another type of server computer[.]” (D.I. 1032-12, Ex. T at VMWCC00016851 (2015/09/23 Appeal Brief at 10).)

Cirba's lone prosecution history excerpt does not require its additional language. Although the applicant said that the "cloud-connector server" "interfaces to users on a remote computer and [] interfaces to remote cloud-connector nodes" (*id.* at VMWCC00016852 (2015/09/23 Appeal Brief at 11)), it was merely paraphrasing existing the claim language without limiting it further. Claims 10 and 19 already recite that the "cloud-connector server" interfaces with remote computers (cloud-computing facilities) and cloud connector nodes. They state that the "cloud-connector server" provides "an electronic cloud-aggregation management interface within one of the multiple, operationally distinct cloud-computing facilities" and "accesses services provided by the cloud-connector node" through "an electronic interface." The applicant's statement thus was neither definitional nor a clear and unmistakable surrender of claim scope. *See Omega Eng'g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1325-26 (Fed. Cir. 2003) (courts have "consistently rejected prosecution statements too vague or ambiguous to qualify as a disavowal of claim scope.").

d. Densify's Sur-Reply Position

During prosecution, the patentee provided an express definition of this term and used it to overcome a prior art rejection, arguing that "[n]othing in paragraph [0045] [of the Martin reference] even remotely suggests that the virtual machines launched by virtual-data-center users provide an interface to a cloud-connector server or access a cloud-management interface within the cloud-computing facility in which they are installed." Ex. B-4 at 15. The definition could not have been more explicit:

The cloud-connector server is thus a dedicated server computer or virtual machine running within a server computer that interfaces to users on a remote computer and that interfaces to remote cloud-connector nodes.

Id. at 11. VMware wants to ignore this express definition and replace it with the “well-known” definition of “server.” *Supra* § III.C.2.c at page 123. The law on this is unambiguous and dispositive: *CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1366 (Fed. Cir. 2002) (“[T]he claim term will not receive its ordinary meaning if the patentee acted as his own lexicographer and clearly set forth a definition of the disputed claim term in either the specification or prosecution history.”). VMware’s arguments to the contrary fail as a matter of law.

3. “cloud-connector node” (’638 Patent: claims 10 and 19)

VMware’s Proposed Construction	Densify’s Proposed Construction
“an application or set of applications that connects the cloud-computing server with a cloud-computing facility”	“an application or set of applications running within a virtual machine installed within each computing-facility of the multiple cloud-computing-facility aggregation, in a virtual data center management server, cloud director, or management system”

a. VMware’s Opening Position

“Cloud-connector node” refers to “an application or set of applications that connects the cloud-computing server with a cloud-computing facility.” It is not indefinite.

Generally, a “node” is well-known in the art as a point of connection in a network. (*See* Ex. A-4 (*Dictionary of Computing* (6th ed. 2010)) at 232 (defining “node” as “an interconnection point in a structure or network”); Ex. A-3 (*Newton’s Telecom Dictionary* (25th Anniversary ed. 2009)) at 782 (defining “node” as “[a] point of connection into a network”).) The intrinsic evidence indicates that the ’638 patent’s “cloud-connector node” is an application or set of applications that connects the cloud-computing server with a cloud-computing facility.

Claims 10 and 19 require that the “cloud-connector node: (i) provide “an electronic interface through which the cloud-connector server accesses services provided by the cloud-connector node” and (ii) “access[] a cloud-management interface within the cloud-computer

facility in which the cloud-connector node is installed” to access “local cloud-management services on behalf of the cloud-connector server.” Dependent claims 14 and 23 limit “each cloud-connector node [to be] a virtual appliance.”

The ’638 patent specification explains that the “cloud-connector node” that meets these requirements is as an application or set of applications that connects the cloud-computing server with a cloud-computing facility.²⁵ It explains that a “cloud connector node” allows the “cloud-connector server” to access services provided by the cloud-computing facility in which the “cloud connector node” is located. (’638 patent, 12:51-55, 13:15-36, FIG. 10.) The “cloud-connector nodes” cooperate with the “cloud-connector server” to provide services that are distributed across multiple cloud-computing facilities. (*Id.*, 12:33-38.)

The specification describes an example in which the “cloud-connector node” is “a web service that executes within an Apache/Tomcat container that runs as a virtual appliance within a cloud director, VDC management server, or third party cloud-computing server.” (*Id.*, 13:37-41; *see also id.*, 12:51-61 (“VCC server 1014 . . . interfaces to a VCC node 1016 that runs as a virtual appliance within the cloud director 1010.”).) In another example, a “cloud-connector node” is “deployed as a virtual appliance, containing one virtual machine in this case, and is launched within the cloud-director servers and VDC management servers[.]” (*Id.*, 14:4-11.)

The prosecution history confirms VMware’s construction of “cloud-connector node.” The patentee explained that “[a] cloud-connector node is generally implemented as an application or set of applications running within a virtual machine.” (D.I. 1032-12, Ex. T at VMWCC00016853 (2015/09/23 Appeal Brief at 12); *see also id.* (“The cloud-connector nodes

²⁵ The specification also references a “cloud-connector node” as a “virtual-cloud-connector node” or a “VCC node.”

are also implemented, in the described implementations, as virtual machines executing within physical server computers.”.)

b. Densify’s Answering Position²⁶

As with the previous term, the patentee admitted to the Patent Office that “cloud-connector node” is not a term known in the art. Ex. B-4 (’638 File History, September 23, 2015 Appeal Brief) at 12 (“Please note that the phrases ‘cloud-connector server’ and ‘cloud-connector node’ do not have well-known or accepted meanings within the computing arts.”).

The patentee also told the Patent Office that a cloud-connector node “is generally implemented as *an application or set of applications running within a virtual machine* (Ex. B-4 (’638 File History, September 23, 2015 Appeal Brief) at 12) and that:

[e]ach of the currently claimed multiple cloud connector nodes is *installed in a virtual data center management server, cloud director, or management system within a cloud-computing facility of the multiple cloud-computing facility*, and each cloud-computing facility of the multiple cloud-computing-facility includes a cloud-connector node.

(D.I. 1032-12, Ex. T at 26 (’638 File History, May 4, 2016 Reply Brief at 7).) This clear lexicography supports Densify’s proposed construction.

For the same reasons expressed above with respect to the “cloud-connector server” term, the patentee’s definition from the prosecution history should control for the “cloud-connector node” term here. *Springs Window Fashions LP v. Novo Indus., L.P.*, 323 F.3d 989, 995 (Fed. Cir. 2003) (“The public notice function of a patent and its prosecution history requires that a patentee be held to what he declares during the prosecution of his patent.”).

²⁶ Densify withdraws its initial construction of “indefiniteness” for this term and proposes the construction herein.

c. VMware's Reply Position

Cirba's attempt to limit "cloud-connector node" to an application or set of applications "running within a virtual machine" is improper for several reasons.

First, Cirba ignores that dependent claims 14 and 23 expressly limit the "cloud-connector node [to be] a virtual appliance," *i.e.*, an application or set of applications running within a virtual machine. The "cloud-connector node" of independent claims 10 and 19 thus is presumed to be broader than claims 14 and 23 under the doctrine of claim differentiation and should not be limited to applications "running within a virtual machine."

Second, Cirba ignores that the specification describes the "cloud-connector node" broadly as an application or set of applications that connects the cloud-computing server with a cloud-computing facility. ('638 patent, 12:33-38 ("cloud-connector nodes" cooperate with cloud-connector server to provide distributed services), 12:51-55, 13:15-36 ("cloud connector node" allows "cloud-connector server" to access services by cloud-computing facility); FIG. 10.)

Third, Cirba tries to make too much of a single statement in the prosecution history, *i.e.*, that "[a] cloud-connector node is *generally* implemented as an application or set of applications running within a virtual machine." (*Supra* § III.C.3.b at page 127 (quoting D.I. 1032-12, Ex. T at VMWCC00016853 (2015/09/23 Appeal Brief at 12)) (emphasis added).) Cirba overlooks that the applicant's statement provided an *example* of a "cloud connector node," but did not define the phrase as limited to applications "running within a virtual machine." Cirba also overlooks the applicant's express use of the word "generally" indicates that other examples are possible.

Nothing reveals a clear and unmistakable disclaimer of claim scope during prosecution. *See Hill-Rom*, 755 F.3d at 1371-72 (disclosure of one embodiment insufficient to limit scope absent manifest exclusion or restriction); *Martek Biosciences Corp. v. Nutrinova, Inc.*, 579 F.3d

1363, 1377 (Fed. Cir. 2009) (“no clear and unmistakable disavowal” where “selected statements” from prosecution were “undercut considerably by additional statements [during] prosecution”).

Cirba’s proposed construction also improperly requires that the cloud-connector nodes be “installed within each computing facility of the multiple cloud-computing-facility aggregation, in a virtual data center management server, cloud director, or management system.” Again relying solely on the prosecution history, Cirba points to the applicant’s statement that “[e]ach of the currently claimed multiple cloud connector nodes is installed in a virtual data center management server, cloud director, or management system within a cloud-computing facility of the multiple cloud-computing facility.” (*Supra* § III.C.3.b at page 127.)

Cirba overlooks that this statement neither limits nor defines “cloud-connector node.” Rather, it refers to structural language *already* recited in the claims: “deploying, registering, and configuring a multiple cloud-connector node in a virtual data center management server, cloud director, or management system within each virtual data center and organization virtual data center in each cloud-computing facility.” (’638 patent, claims 10, 19.) Cirba’s additional language thus would render express claim language superfluous.

d. Densify’s Sur-Reply Position

For this term too, the patentee provided an express definition and relied on it to overcome a prior art rejection. Specifically, the patentee stated that a cloud-connector node “is generally implemented as *an application or set of applications running within a virtual machine*” and that “[e]ach of the currently claimed multiple cloud connector nodes is *installed in a virtual data center management server, cloud director, or management system within a cloud-computing facility of the multiple cloud-computing facility*, and each cloud-computing facility of the multiple cloud-computing-facility includes a cloud-connector node.” *See Supra* § III.C.3.b at page 127.

VMware first criticizes inclusion of the “running within a virtual machine” language, arguing that the patentee’s statement of how the cloud-connector node is “generally implemented” broadens the scope of the claim term. But this ignores the patentee’s statement on the very same page that the term did “not have well-known or accepted meanings within the computing arts” and that the “phrases must be interpreted *according to the definitions and explanations provided in the current application.*” Ex. B-4 at 12 (emphasis added). In the absence of any other definitions or explanations, a POSA would have no interpretation of the claim term outside of the record definition. And VMware’s claim differentiation argument should be rejected where, as here, the claim differentiation “presumption can be overcome by lexicography or disavowal of claim scope.” *See, e.g., Shire Dev. LLC v. Teva Pharms. USA, Inc.*, 2019 WL 969638, at *8 (D. Del. Feb. 28, 2019).

VMware next takes issue with the requirement that the cloud-connector nodes be “installed within each computing facility of the multiple cloud-computing facility aggregation, in a virtual data center management server, cloud director, or management system.” *Supra* § III.C.3.c at page 129. But again, this tracks patentee’s own definition, and VMware does not explain what else would have guided a POSA’s understanding of the bespoke terminology. VMware’s disregard of the only record evidence defining the term and its efforts to broaden the claim in spite of that evidence should be rejected.

D. U.S. Patent No. 10,261,842 (“’842 Patent”)

1. “distributed resource scheduling module” (’842 Patent: claim 1)

VMware’s Proposed Construction	Densify’s Proposed Construction
Not a means-plus-function term subject to §112, ¶ 6 and not indefinite. This term should not be construed as its takes its plain and ordinary meaning.	This is a means-plus-function term subject to § 112, ¶ 6. The term is indefinite for lack of sufficiently definite corresponding structure.

a. VMware's Opening Position

The '842 patent (D.I. 1032-4, Ex. N) describes novel methods and systems to account for expected future workloads in a virtual computing environment. The inventions accommodate the effects of implementing a recommended change (*e.g.*, migrating a VM to a new host) on both the current workloads and future workloads in the computing environment. (*Id.*, 1:47-51, 2:54-61, 9:15-17, 9:27-10:11, 10:17-19, 11:64-12:11, FIG. 5.) If a combined effect on current and future workloads (*e.g.*, a metric of improvement in load-balancing in the environment) exceeds a pre-defined threshold, the disclosed methods perform the recommended change. (*Id.*)

“Distributed resource scheduling module” is not a means-plus-function limitation. Cirba cannot “overcome the presumption that § 112, ¶ 6 does not apply to claim terms not including the word ‘means.’” *M2M Sols. LLC v. Sierra Wireless Am., Inc.*, No. 12-30-RGA, 2016 WL 1298961, at *2, 5-6 (D. Del. Mar. 31, 2016) (“memory module” and “processing module” not means-plus-function terms); *see also Zeroclick, LLC v. Apple Inc.*, 891 F.3d 1003, 1008 (Fed. Cir. 2018) (“[T]hat the disputed limitations incorporate functional language does not automatically convert the words into means for performing such functions.”). As the phrase “distributed resource scheduling module” has “a known structural meaning, or recites either a known or generic term with sufficient description of its operation,” “the presumption against means-plus-function claiming remains intact.” *M2M Sols. LLC v. Sierra Wireless Am., Inc.*, No. 12-30-RGA, 2015 WL 5826816, at *3-4 (D. Del. Oct. 2, 2015) (emphasis omitted).

Distributed resource scheduling modules were well-known in the art. Cirba's own expert Dr. Madisetti established this earlier in the case, when opining about the similar phrase “resource allocation module” in a different VMware patent, the '752 patent. Dr. Madisetti declared that “a POSA would have been generally familiar with the concept of resource allocation.” (D.I. 789-2, Ex. B-1 ¶ 25.) And he conceded that “VMware's distributed resource scheduler (DRS) could

have been considered a resource allocation module, [and] DRS was far from the only product on the market that performed that function. Some of other products on the market at that time include Citrix NetScaler and Microsoft Virtual Machine Manager Hyper-V.” (*Id.* ¶ 30.)

Dr. Madisetti thus recognized that distributed resource schedulers, and their resource allocation function, were well-known in the art. (*Id.*) His view is consistent with VMware’s expert, Dr. Menascé. When discussing the ’752 patent, Dr. Menascé too declared that distributed resource schedulers were well known. (D.I. 789-1, Ex. A-1 ¶¶ 10-13; *id.*, Ex. A-2 ¶¶ 35-36.) Cirba’s current position—that “distributed resource scheduling module” is a means-plus-function limitation—thus conflicts with both experts’ earlier position to the contrary.

The background section of the ’842 patent’s specification identifies VMware’s Distributed Resource Scheduler as an example of a distributed resource scheduling module. (’842 patent, 1:28-32.) The patent then discloses that “[c]onventional techniques for Distributed Resource Scheduling (DRS) and Distributed Power Management (DPM) operate in a reactive mode to demand changes, where VM migration and host power-ons and power-offs are recommended ‘reactively’ based on current VM demand data.” (*Id.*, 1:47-51.) Section 112, ¶ 6, does not apply to such well-known claim elements. *M2M Sols. LLC*, 2016 WL 1298961, at *2, 5-6; *see also Citrix Sys., Inc. v. Workspot, Inc.*, No. 18-588- 2020 WL 5634219, at *9-10 (D. Del. Sept. 21, 2020) (“network module for receiving . . .”, “authentication module for authenticating . . .”, and “server process for reestablishing . . .” were “generally known in the art” and not subject to § 112, ¶ 6).

The ’842 patent discloses and claims novel inventions for improving systems and methods that employ existing, well-known distributed resource scheduling modules. The patent

enables these components to recommend changes to a virtual computing environment based on both current and future workloads. (*Id.*, claim 1; *see also id.*, 2:54-61, 9:27-10:11, FIG. 5.)

Because a distributed resource scheduling module does not use the word “means” and was well-known in the art, the term is not subject to § 112, ¶ 6 and the Court need not construe it. Alternatively, the Court could construe “distributed resource scheduling module” to take its plain and ordinary meaning as a “distributed resource scheduler.”

b. Densify’s Answering Position

The term “distributed resource scheduling module” is subject to 35 U.S.C. § 112, ¶ 6 and is indefinite because the specification does not disclose adequate corresponding structure.

A means-plus-function inquiry has two parts. First, the Court must determine whether the claim limitation employs means-plus-function format. Where, as here, a claim does not use the word “means,” there is a rebuttable presumption that means-plus-function claiming does not apply. The presumption may be overcome by showing by a preponderance of the evidence, that “the claim term fails to recite sufficiently definite structure.” *Zeroclick, LLC v. Apple Inc.*, 891 F.3d 1003, 1007 (Fed. Cir. 2018); *Williamson v. Citrix Online, LLC*, 792 F.3d 1339, 1349 (Fed. Cir. 2015). The question is “whether the words of the claim are understood by [POSAs] to have a sufficiently definite meaning as the name for structure.” *Zeroclick*, 891 F.3d at 1007. Once a claim is determined to be “means-plus-function,” the court then identifies the function and assesses whether the specification sufficiently discloses a corresponding structure. *Williamson*, 792 F.3d at 1351.

Here, means-plus-function claiming applies because the term “distributed resource scheduling module” is a nonce term. *Williamson*, 792 F.3d at 1350 (“‘Module’ is a well-known nonce word that can operate as a substitute for ‘means’ [It] is simply a generic description for software or hardware that performs a specified function.”); *see also* M.P.E.P § 2181 (“module,”

like “mechanism for,” “device for,” or “system for,” is “a non-structural generic placeholder[]); *TQ Delta, LLC v. 2Wire, Inc.*, 2018 WL 626472, at *9-10 (D. Del. Jan. 30, 2018) (“message determination module” subject to means-plus-function claiming). Simply appending an operation (“distributed resource scheduling”) to the term “module” does not provide any information as to what the relevant structure is. A POSA may have been familiar with “resource allocation” (or even “resource scheduling”) generally (*see Supra* § III.D.1.a at page 132), but that does not change the fact that the *mechanism* for carrying out the functions is subjective. Madisetti Decl., dated June 9, 2020 at ¶¶ 25 – 27, 29 (D.I. 789-2, Ex. B-1); Ex. B-2 at ¶¶ 37-44. There are many ways it could have been done. *Id.*

(i) **Functions Carried Out By The Claimed “Module”
Include “Receiving..., Determining..., Calculating...,
and Performing...”**

Under applicable law, the disclosed structure of the claimed “module” must perform all functions associate with it. *Noah Sys., Inc. v. Intuit Inc.*, 675 F.3d 1302, 1318-19 (Fed. Cir. 2012) (“[W]here a disclosed algorithm supports some, but not all, of the functions associated with a means-plus-function limitation, we treat the specification as if no algorithm has been disclosed at all.”). Here, those functions are:

- ***receiving*** a recommended change to a virtual architecture of the virtual computing environment ***at a distributed resource scheduling module***,...
- ***determining, by the distributed resource scheduling module***, an impact on current workload in the virtual computing environment if the recommended change is performed;
- ***determining, by the distributed resource scheduling module***, an impact on future workload in the virtual computing environment if the recommended change is performed;
- ***calculating, by the distributed resource scheduling module***, a combined impact on current and future workload from the determined impact on current workload in the virtual computing environment if the recommended change is performed and from the determined impact on future workload in the virtual computing environment if the recommended change is performed;
- ***determining, by the distributed resource scheduling module***, that the combined impact is above a threshold; and

- in response to determining that the combined impact on current and future workload is above the threshold, **performing** the recommended change to the virtual architecture of the virtual computing environment.

'842 patent, claim 1. Each of the “receiving,” “calculating” and “determining” functions are, by the plain language of the claim, associated with the module. The last element requiring “performing the recommended change...” does not expressly contain the “distributed resource scheduling module” language, but based on the structure of the claim as a whole, that module must be responsible for carrying out that function; no other structure is recited or associated with it.

(ii) The Specification Does Not Disclose Corresponding Structure.

The structure associated with each of the claimed functions is the “distributed resource scheduling module.” Nothing in the language of the claim provides information regarding what comprises the module or how it carries out the claimed functions. It is a nonce phrase – a generic description for something that performs specified functions. *Williamson*, 792 F.3d at 1350.

The '842 patent specification provides no adequate corresponding structure for the “distributed resource scheduling module.” The specification describes the functions that the “DRS module” performs and further provides a diagram of the functions at Figure 5. '842 patent at Fig. 5; 9:27-10:11. But both the descriptions and Figure 5 fail to provide sufficient information regarding the structure of the claimed “module” or any particular algorithm(s) for performing its associated functions. At most, the specification mentions a “DRS algorithm” one time, but does not provide or describe the algorithm or any other relevant detail. '842 patent at 6:32-36. That is not sufficient under § 112, ¶ 6. Indeed, a POSA would have understood that multiple possible algorithms or techniques could be used or designed based on a variety of factors. Ex. B-2 at

¶¶ 37-44; *Arendi S.A.R.L. v. LG Elecs., Inc.*, 2019 WL 3891150, at *13 (D. Del. Aug. 8, 2019) (J. Stark) (holding “the multiplicity of possible algorithms ... and the specification’s lack of any disclosure of any particular algorithm ..., [causes] the Court [to] conclude[] that the specification fails to disclose sufficient structure for the [relevant] function and, thus, the limitation is indefinite”).

VMware relies on the patent’s mention of “[c]onventional techniques for Distributed Resource Scheduling (DRS) and Distributed Power Management (DPM) operate in a reactive mode to demand changes, where VM migration and host power-ons and power-offs are recommended ‘reactively’ based on current VM demand data” (‘842 patent, 1:47-51) and argues that “Section 112, ¶ 6, does not apply to such well-known claim elements.” *Supra* § III.D.1.a at page 132. But there is no discussion of what these allegedly “conventional techniques” are. *Cirba Inc. v. VMware, Inc.*, No. CV 19-742-LPS, 2020 WL 5910112, at *13 (D. Del. Oct. 6, 2020) (J. Stark) (“The issue is not whether a POSA would know, given the specification, how to implement the claimed ... function, but instead whether a POSA would recognize *the specification itself as disclosing a particular algorithm or algorithms for implementing the ... function.*”) (emphasis added).

VMware’s reliance on *M2M Solutions LLC v. Sierra Wireless America, Inc.*, 2016 WL 1298961 (D. Del. Mar. 31, 2016) is inapposite. The *M2M* Court found that means-plus function claiming did not apply because the claim included additional language that “describes how the authentication process takes place in considerable detail.” *M2M*, 12-cv-30, D.I. 215 at 8-9 (D. Del. Oct. 2, 2015) (“Despite the fact that the claim recites a function, the immediately following words provide algorithmic structure for performing that function.”). The claim here does not describe how the distributed resource scheduling module works. Moreover, *M2M* found that the

defendant failed to meet its burden by not submitting any evidence of how a POSA would view the claim term and not addressing the claim language describing how the claimed function was carried out. *Id.* at 8; *see also M2M*, 2016 WL 1298961, at *6 (“Defendants simply offer conclusory attorney statements....”).

c. VMware’s Reply Position

(i) Distributed Resource Scheduling Modules are Structures Known in the Art and Thus Not a Means Plus Function Claim Term

“Distributed resource scheduling module” (“DRS module”) is not subject to 35 U.S.C. § 112, ¶ 6. It is undisputed that “DRS module” refers to a specific class of structures known in the art. A POSITA understands that “DRS module,” as of the ’842 patent’s filing date, referred to a conventional class of software programs that perform management of distributed resources, such as VM workloads. (Ex. A-20 ¶¶ 22-25; Ex. B-2 ¶ 42 (“A POSA would have also understood that multiple commercially-available products were available that utilize techniques for resource scheduling, such as Citrix Workload Balancing (WLB), Microsoft Virtual Machine Manager Hyper-V Hypervisor Scheduler, [and] Red Hat Virtualization Cluster Scheduler.”).)

This use of a conventional class of structures without the term “means” does not invoke § 112, ¶ 6. *See Personalized Media Commc’ns, LLC v. ITC*, 161 F.3d 696, 705 (Fed. Cir. 1998) (“Even though the term ‘detector’ does not specifically evoke a particular structure, it does convey to one knowledgeable in the art a variety of structures known as ‘detectors.’ We therefore conclude that the term ‘detector’ is a sufficiently definite structural term to preclude the application of § 112, ¶ 6.”); *Linear Tech. Corp. v. Impala Linear Corp.*, 379 F.3d 1311, 1322 (Fed. Cir. 2004) (“That the disputed term is not limited to a single structure does not disqualify it as a corresponding structure, as long as the class of structures is identifiable by a [POSITA].”).

The intrinsic evidence confirms the conventional nature of a DRS module. As VMware explained in its opening brief, the specification teaches that well-known DRS modules include VMware’s Distributed Resource Scheduler and Distributed Power Management software. (’842 patent, 1:47-56.) Cirba does not rebut this. During prosecution, the applicant also explained that the claimed method “focuse[s] on a specific asserted technical improvement in computer capabilities” (*i.e.*, calculating the combined current and future impacts of recommended changes) over “conventional techniques for *Distributed Resource Scheduling (DRS)* and Distributed Power Management (DPM) in a virtual computing environment.” (D.I. 1032-14, Ex. U at VMWCC00017304, -06 (2018/10/9 Amend. And Resp. at 10, 12) (emphasis added).)

Besides attorney argument, Cirba offers no evidence that the term “DRS module” is not structural and thus is subject to § 112, ¶ 6. (*Supra* § III.D.1.b at pages 136-37.) Its expert, Dr. Madisetti, intentionally avoided offering an opinion on the topic. He stated he had been “asked by counsel for Densify . . . to *assume* that ‘distributed resource scheduling module’ is a means-plus-function term subject to interpretation under § 112, ¶ 6.” (Ex. B-2 ¶ 37 (emphasis added).) Dr. Madisetti’s evasion is unsurprising, as he previously opined that DRS modules like VMware’s DRS and Citrix Workload Balancing were well-known in the art and commercially-available. (*Supra* § III.D.1.a at page 134-35; D.I. 789-2, Ex. B-1 (June 9, 2020 Madisetti Decl.) ¶¶ 25-26, 28-30.)

Absent factual or expert support for classifying “DRS module” as non-structural, Cirba has not overcome the presumption that § 112, ¶ 6, is inapplicable as a matter of law.

**(ii) Even if “Distributed Resource Scheduling Module”
Were Subject to § 112, ¶ 6, the ’842 Patent Discloses
Sufficient Structure**

Even if Cirba were correct that “DRS module” is subject to § 112, ¶ 6, the patent sufficiently describes structures that carry out the recited functions, *viz.*, “receiving a

recommended change,” “determining ... an impact on current workload,” “determining ... an impact on future workload,” “calculating ... a combined impact,” and “determining ... that the combined impact is above a threshold.”

As an initial matter, Cirba focuses too narrowly on Figure 5 in alleging that this figure and its related descriptions “fail to provide sufficient information regarding the structure of the claimed ‘module’ or any particularly algorithm(s) for performing its associated functions.” (*Supra* § III.D.1.b at pages 138-39.) Cirba ignores that the claims themselves and the specification as a whole describe a specific algorithm and details for the recited steps. (Ex. A-20 ¶¶ 28-37)

Per the plain language of claim 1, the DRS module “receiv[es] a recommended change to a virtual architecture of the virtual computing environment,” which “includes at least one of migrating a virtual machine . . . , a computer power-on, a host computer power-off, and a migration of a disk image.” The module also “determin[es] ... [the] impact on current workload” and “on future workload in the virtual computing environment” from performing the recommended change. The module further “calculat[es] ... a combined impact on current and future workload” and “determin[es] ... [whether] the combined impact is above a threshold.”

The POSITA thus understands that the claim itself sets forth a specific algorithm for carrying out the claimed method steps associated with the DRS module. (Ex. A-20 ¶ 28; *see also* Ex. B-2 ¶ 43 (describing a POSA’s considerations when designing an algorithm).) The claim language specifies the inputs, outputs, and objectives of the claimed algorithm. This is legally sufficient structure. *See Align Tech., Inc. v. 3Shape*, No. 17-cv-1648-LPS, 2021 WL 2320139, at *7 (D. Del. June 7, 2021) (claim disclosed sufficient structure where it “provide[d] an input-

output structure” for the claim element and “explain[ed] how the [element] interacts with other components of the claim.”) (internal quotation and citation omitted).

Besides the plain claim language, the specification provides details for each step:

“receiving . . . a recommended change.” The specification teaches that recommended changes are “received” by the DRS module from a VM management computer “configured to transmit the recommendation for resource allocation to the DRS module.” (’842 patent, 8:61-66; *see also id.*, 9:36-40, Fig. 5.) The DRS module is implemented within the VM management computer, which is interconnected to hosts and VMs to “view and manage resources in the virtual computing environment.” (*Id.*, 5:24-33, Fig. 2A; Ex. A-20 ¶¶ 30-31.)

A POSITA understands that this structure calls for conventional software modules that receive recommendations for evaluation, such as the commercially available ones that Dr. Madisetti describes. (Ex. B-2 ¶ 42; Ex. A-20 ¶¶ 24-25, 30-31.) Confirming this, the specification informs a POSITA that the “recommended change” to the virtual architecture “receiv[ed]” by the DRS module includes conventional forms of distributed resource scheduling in a virtual environment, such as “launching VM migrations and/or host power-ons and power-offs.” (’842 patent, 1:47-56; *see also id.*, 5:34-44 (“In some embodiments, the DRS module . . . makes automatic resource relocation and optimization decisions as hosts or virtual machines are added or removed from the virtual environment.”), Fig. 2A, Fig. 5; Ex. A-20 ¶ 31.)

“determining an impact on the current workload.” The specification explains that the DRS module “determines an impact on the current workload if the recommended change is performed,” *e.g.*, by scoring VM migrations “based on [their] improvement in the cluster imbalance metric and on its risk-adjusted costs and benefits” before the “imbalance metric is recomputed incorporating that move.” (’842 patent, 9:1-26; *see also id.*, 9:41-45, Fig. 5.) Like

the recommended change, this is a conventional determination when performing resource allocation and distributed resource scheduling. (Ex. A-20 ¶¶ 32.)

“determining an impact on the future workload.” Just as conventional DRS modules determine an impact on the current workload using current statistics, the patent explains that the disclosed DRS module “determines an impact on the future workload if the recommended change is performed” using a “predicted load.” (’842 patent, 7:42-8:7; *see also id.*, 9:45-49, Fig. 5.) The specification discloses several exemplary methodologies for determining the future workload, including a “pattern matching predictor, a “‘polyfit’ prediction mechanism,” and a “‘simple step prediction’ mechanism.” (*Id.*, 8:13-60.) It also further details particular exemplary approaches to “pattern matching” (*id.*, 6:11-8:40) and “simple step prediction” (*id.*, 8:41-60). The POSITA understands that future workloads also could be specified in other ways (*e.g.*, user pre-defined future states). (Ex. A-20 ¶¶ 33-34.)

“calculating ... a combined impact on current and future workload.” As noted above, claim 1 itself identifies inputs (*i.e.*, the determined impacts on current and future workload) to the DRS module that are used to calculate an output (*i.e.*, a combined impact). The specification also identifies exemplary methods for calculating the combined impact, such as “scor[ing] as the sum of current and future impact, weighted by confidence in the future prediction.” (’842 patent, 9:50-67.) These examples guide the POSITA regarding an appropriate methodology for calculating combined impacts. (Ex. A-20 ¶¶ 35-36; *see also* Ex. B-2 ¶ 43 (POSITA would understand how to design an appropriate algorithm).)

“determining ... that the combined impact is above a threshold.” Per the specification, the DRS module compares the combined impact to a threshold to see if a recommended change should be performed, *e.g.*, as part of a cost-benefit analysis. (’842 patent, 9:50-10:11; *see also*

id., 8:1-7 (“Some embodiments . . . include performing a cost-benefit analysis to determine whether the recommendations suggested by the VM management center should be performed....”), 9:27-30 (“FIG. 5 is a flow diagram of method steps for performing a cost-benefit analysis to determine whether [suggested] recommendations . . . should be performed.”).) A POSITA understands that many thresholds could be compared against the combined impact, such as a “user-configurable threshold.” (*Id.*, 9:12-15; Ex. A-20 ¶ 37.)

Thus, even assuming *arguendo* that “DRS module” were subject to § 112, ¶ 6 (which it is not), the ’842 patent still would disclose sufficient corresponding structure for the recited functions, and the claims would not be indefinite. (Ex. A-20 ¶¶ 38-40.)

d. Densify’s Sur-Reply Position

Module is a common nonce word subject to 35 U.S.C. §112, ¶ 6. VMware’s reliance on *Personalized Media Commc’ns, LLC v. Int’l Trade Comm’n*, 161 F.3d 696, 704 (Fed. Cir. 1998) for “detector” is misplaced. The *Personalized* court concluded that “detector” was not a generic structure term such as “means,” “element,” or “device,” or in this case “module.” Instead, the court noted that, according to dictionary definitions, “detector” had a well-known meaning to those of skill in the electrical arts connotative of structure, “including a rectifier or demodulator.” *Id.* at 705. Here, “module” is a classic nonce word that does not connote structure. VMware has not shown any references to a class of structures for a DRS module. VMware’s criticism of Dr. Madisetti’s declaration is inapposite – the declaration was discussing resource scheduling *techniques*, and not *structures*, for a *different* patent. Ex. B-2 ¶¶ 37-44. Dr. Dinda’s declaration, citing the claim language and portions of the written description, fails to describe *how* the distributed resource scheduling module, by its interaction with the other components in the scheduling distributed computer systems, may serve as a structure. And VMware’s Distributed Resource Scheduling (DRS) and Distributed Power Management (DPM) are both

proprietary black boxes, which have no publicly-known structure. Technical dictionaries do not define distributed resource scheduling modules to those skilled in the computer arts, nor does the prefix “distributed resource scheduling” impart the required structure into the term “module.” These words do not describe a sufficiently definite structure. Ex. B-2 ¶¶ 37-44.

The '842 Patent Provides No Algorithm for the DRS Module. VMware does not show how the '842 patent provides an algorithm in the specification and the figures. Ex. B-2 ¶ 39-43. VMware contends that the specification describes the inputs, outputs, and objectives of the algorithm by restating the claimed functions, but the claimed functions are not the algorithm, but simply describe the function to be performed. *Aristocrat Techs. Australia Pty Ltd. v. Int'l Game Tech.*, 521 F.3d 1328, 1334 (Fed. Cir. 2008) (“But that language simply describes the function to be performed, not the algorithm by which it is performed.”); *see also Finisar Corp. v. DirecTV Grp., Inc.*, 523 F.3d 1323, 1340 (Fed. Cir. 2008) (“this passage provides ‘nothing more than a restatement of the function, as recited in the claim.’”) VMware’s argument merely parrots the claimed functions in the specification and the figures (*See* Fig. 5). Dr. Dinda’s declaration does not explain how any algorithm is described within the specification, only that one skilled in the art would know how to design particular software modules for carrying out this claim limitation. That’s not enough to overcome indefiniteness. As this Court found in *Arendi*, the issue “is not whether a POSA would know, given the specification, how to implement the claimed...function, but instead whether a POSA would recognize the *specification itself* as disclosing a particular algorithm or algorithms for implementing the...function.” *Arendi S.A.R.L. v. LG Elecs., Inc.*, 2019 WL 3891150, at *13; *see also Function Media, L.L.C. v. Google, Inc.*, 708 F.3d 1310, 1319 (Fed. Cir. 2013) (“[I]t is well established that proving that a person of ordinary skill could devise

some method to perform the function is not the proper inquiry as to definiteness—that inquiry goes to enablement”).

2. **“A non-transitory computer-readable storage medium comprising instructions that, when executed in a computing device, causes the computing device to carry out the steps of” (claim 12) (’842 Patent: claim 12)²⁷**

VMware’s Proposed Construction	Densify’s Proposed Construction
Not a means-plus-function term subject to §112, ¶ 6 and not indefinite. As preambles, these terms are not limiting, merely identify well-known claim forms (non-transitory computer-readable storage media and computer systems), and should not be construed.	This is a means-plus-function term subject to § 112, ¶ 6. The term is indefinite for lack of sufficiently definite corresponding structure.

a. VMware’s Opening Position

These phrases constitute the preambles to independent claims 12 and 17. Because they do not use the word “means” and are well-known preamble components used in their conventional sense, Cirba cannot overcome the presumption against § 112, ¶ 6. *See Zeroclick, LLC*, 891 F.3d at 1008; *M2M Sols. LLC*, 2016 WL 1298961, at *2, 5-6; *M2M Sols. LLC*, 2015 WL 5826816, at *3-4. None of these phrases even recites a *function* that might invoke § 112, ¶ 6, which forecloses treating them as means-plus-function limitations.

The phrases “non-transitory computer-readable storage medium” and “computer system” are commonplace components that are routinely invoked as non-limiting claim preambles. They are widely accepted to describe the form of a claimed invention in the computer arts satisfying 35 U.S.C. § 101. *See* MPEP 2111.05. Sec. III (describing evaluating scope of claim to computer-readable medium); D. Kappos, “Subject Matter Eligibility of Computer Readable

²⁷ The parties originally grouped claims 12 and 17 (“A computer system, comprising: one or more host computers; and a virtual management computer, configured to” (claim 17)) together, but the terms were subsequently addressed separately by Densify in its Answering Brief.

Media,” 1351 Off. Gaz. Pat. Office 212 (Feb. 23, 2010) (Director of the USPTO stating that the “USPTO suggest[s]” adding “non-transitory” to the preamble of “computer readable medium” claims to avoid or overcome rejections under 35 U.S.C. § 101); *cf. Symantec Corp. v. Computer Assocs. Int’l, Inc.*, 522 F.3d 1279, 1291-92 (Fed. Cir. 2008) (construing “computer system” to have its plain and ordinary meaning in the body of a claim). These preamble phrases are not limiting, are not means-plus-function limitations, and are not indefinite.

b. Densify’s Answering Position²⁸

Claim 12 reads:

A non-transitory computer-readable storage medium comprising instructions that, when executed in a computing device, ***causes the computing device to carry out the steps of:***

receiving a recommended change to a virtual architecture of a virtual computing environment, wherein the recommended change includes at least one of migrating a virtual machine from one host computer to another host computer, a computer power-on, a host computer power-off, and a migration of a disk image;

determining an impact on current workload in the virtual computing environment if the recommended change is performed;

determining an impact on future workload in the virtual computing environment if the recommended change is performed;

calculating a combined impact on current and future workload from the determined impact on current workload in the virtual computing environment if the recommended change is performed and from the determined impact on future workload in the virtual computing environment if the recommended change is performed;

determining if the combined impact is above or below a threshold and selecting between one of the available options of;

if the combined impact on current and future workload is below the threshold, ***do not perform*** the recommended change to the virtual architecture of the virtual computing environment; and

²⁸ Terms from claim 17 and claim 12 were grouped together in the Joint Claim Construction Chart (D.I. 1032), but they are discussed separately here for better readability.

if the combined impact on current and future workload is above the threshold, *perform* the recommended change to the virtual architecture of the virtual computing environment.

(i) The Preamble Is Limiting

The preamble of claim 12 is limiting. “To be limiting, a preamble must affect a manipulative difference in the claims, provide antecedent basis for a claim term, or breathe life into the claims.” *Novartis Pharm. Corp. v. Accord Healthcare Inc.*, 387 F. Supp. 3d 429, 436 (D. Del. 2019) (collecting cases). For claim 12, the preamble recites a necessary element of the invention – without the preamble’s recitation of “computer-readable storage medium comprising instructions,” the claim would have no structure that enables the steps set forth in the body of the claim. In other words, because the patentee does not “define[] a structurally complete invention in the claim body,” the preamble is necessarily limiting. *Fresenius Med. Care Holding, Inc. v. Baxter Int’l, Inc.*, No. C 03-1431 SBA, 2006 WL 1343648, at * 16 (N.D. Cal. May 16, 2006) (quoting *Rowe v. Dror*, 112 F.3d 473, 478 (Fed. Cir. 1997)).

(ii) Claim 12 Is Subject to 35 U.S.C. § 112, ¶ 6

Although claim 12’s lack of the word “means” gives rise to a presumption that 35 U.S.C. § 112, ¶ 6 does not apply, that presumption is overcome because “computer-readable storage medium comprising instructions” is a generic black box for performing the claimed computer implemented steps. As noted, the only structure recited for *all* of the claimed steps is the “computer-readable storage medium comprising instructions,” which is insufficiently detailed and would not have provided a POSA an adequate indication of the structure claimed.

Williamson, 792 F.3d at 1350.

The facts here are identical in relevant part to those before this Court in *Arendi*, 2019 WL 3891150 at *13. In *Arendi*, the Court found claim 98 of the U.S. Patent No. 7,496,854 to be subject to Section 112, ¶ 6. Claim 98 claimed:

A computer readable medium for assisting a computer operator to retrieve information from a database that is related to text in a document, the computer readable medium including program instructions for performing the steps of:

- (1) using a first computer program to analyze the document, without direction from the operator, to identify text in the document that can be used to search for related information,
- (2) using a second computer program and the text identified in step (1) to search the database and to locate related information, and
- (3) inserting the information located in step (2) into the document.

The Court found that even though the claim “does not use the word ‘means,’” it is nevertheless subject to Section 112, ¶ 6 “if the limitation merely uses a nonce word to connote a generic ‘black box’ for performing a computer-implemented method.” *Arendi*, 2019 WL 3891150 at *12 (quoting *Williamson*, 792 F.3d at 1350). And because “the only structure recited for the function of ‘inserting’ is the ‘computer readable medium program including instructions,’” the Court found it was a “generic black box” that did “not give a POSA a sufficiently definite ‘indication of [the] structure’ (i.e., algorithm) claimed,” thus making it proper to treat the claim as “a means-plus-function limitation.” *Id.*

(iii) The Functions Performed By the “Non-Transitory Computer-Readable Storage Medium Comprising Instructions...” Term Are the Steps Listed In the Claim

Each of the “receiving,” “calculating” and “determining” functions are, by the plain language of the claim, performed by the “computer-readable storage medium comprising instructions.” The last two elements that require performing or not performing steps based on the outcome of a prior step are also carried out by the same structure; no other structure is recited.

VMware asserts that “none of these phrases even recites a *function* that might invoke § 112, ¶ 6.” *Supra* § III.D.2.a at page 145 (emphasis in original). This ignores the plain language

of the claim, which expressly calls for doing something: “determining an impact...” and “calculating a combined impact...” etc.

(iv) The Specification Does Not Disclose Corresponding Structure.

“If the patentee fails to disclose adequate corresponding structure, the claim is indefinite.” *Williamson*, 792 F.3d at 1351-52. The ’842 patent specification does not disclose a sufficient structure for the claimed “non-transitory computer-readable storage medium comprising instructions” In particular, the specification fails to disclose sufficient structure for performing the functions claimed by each of these limitations.

As one example, the “calculating” step requires “calculating a combined impact on current and future workload from the determined impact on current workload” However, “the specification does not explain *how*, algorithmically, this [calculating] is to occur.” *Arendi*, 2019 WL 3891150 at *13. The only discussion of calculating combined impacts relates to Fig.

5:

At step **508**, the DRS module calculates a combined impact on current and future workload. At step **510**, the DRS module determines whether the combined impact on current and future workload is above or below a threshold. In some embodiments, moves that hurt both current and future cluster imbalance are rejected. Moves that improve both current and future cluster imbalance are scored as the sum of current and future improvement, weighted by confidence in the future prediction of VMs on the source and target hosts. *Moves that improve current, but hurt future cluster imbalance, could be scored as the improvement in current imbalance, with the idea that cost-benefit analysis will likely reject this move. Moves that hurt current but improve future cluster imbalance are challenging to analyze; these moves need to be considered in light of how much benefit may be derived in the future (along with the confidence in the future) versus how much impact there is on current cluster imbalance.*

'842 patent at 9:53-67 (emphasis added). The disclosure talks about “scoring” and “weighting” moves in certain instances (when moves improve both current and future cluster imbalance) but never says how. In other instances (when moves hurt both current and future cluster imbalance), the patent admits that analyzing them is “challenging” and states that they “need to be” considered in light of certain circumstances, but offers no insight as to how that might be done. Nor does the specification teach how these moves can (or should) be “combined.” Fig. 5 itself offers no insight – the flowchart is comprised solely of black boxes reciting the functions, but no information as to how to do them. And POSA would have understood that multiple possible algorithms or techniques could be used or designed based on a variety of factors. Ex. B-2 at ¶¶ 45-52. “The specification’s lack of any disclosure of any particular algorithm or algorithms” corresponding to the “non-transitory computer-readable storage medium comprising instructions...” renders the term indefinite. *Arendi*, 2019 WL 3891150 at *13.

c. VMware’s Reply Position

(i) The Claimed Instructions Are the Claimed Structure

The parties agree that claim 12’s preamble is limiting to the extent that it specifies “instructions that, when executed in a computing device, cause[] the computing device to carry out the steps of” the method recited in claim 12. But Cirba takes the novel position that “computer-readable storage medium comprising instructions” (“CRM”) claims are necessarily subject to § 112, ¶ 6. (*See supra* § III.D.2.b at page 149 (alleging that presumption against non-application of § 112, ¶ 6 “is overcome because ‘computer-readable storage medium comprising instructions’ is a generic black box for performing the claimed computer implemented steps”).)

Cirba overlooks the claim language itself. Claim 12 provides a structure, as it “recite[s] the objectives and operations” of the CRM limitation and “provides a description of how the

computer-readable medium is specifically programmed to operate.” *Uniloc 2017 LLC v. Google LLC*, No. 2:18-cv-492-JRG-RSP, 2020 WL 569858, at *15 (E.D. Tex. Feb. 5, 2020).

Cirba has no expert support for its contention that CRM claims necessarily invoke § 112, ¶ 6, despite not using the term “means.” (*Compare supra* § III.D.2.b at pages 149-50 *with supra* § III.D.2.a at page 147-48 (citing USPTO guidance).) Not even Cirba’s expert was willing to make this leap. Rather than opine as to whether the CRM preamble connotes structure, Dr. Madisetti again was “asked by counsel for Densify to *assume* that [the CRM preamble] is a means-plus-function term subject to interpretation under § 112, ¶ 6.” (Ex. B-2 ¶ 45 (emphasis added).)

Lacking expert support, Cirba’s sole justification for its view that CRM claims necessarily invoke § 112, ¶ 6 is this Court’s opinion in *Arendi S.A.R.L. v. LG Electronics*, No. 12-1595-LPS, 2019 WL 3891150 (D. Del. Aug. 19, 2019). There, the Court held an “inserting” instruction to be subject to § 112, ¶ 6 because it failed to provide any guidance as to how or where to “insert” the relevant data. *See id.* at *13 (“[T]he specification includes no specific protocols, no flow-charts, and no other description of how the claimed ‘inserting’ function is to be implemented.”) Here, however, claim 1 sets forth particular computer-implemented instructions that must be carried out and how those instructions interact—*i.e.*, receive a recommended change, determine current and future impacts of the recommended change, calculate a combined impact from the current and future impacts, determine if the combined impact exceeds a threshold, and, if it does, perform the recommended change. (’842 patent, claim 12; *see also* Ex. A-20 ¶¶ 41-43.) And unlike the claim term in *Arendi*, the particular set of instructions here is itself an algorithm that informs a POSITA how the instructions must be

executed, including their objectives and operations, as further described in the specification. (*infra* § III.D.1.c(ii); Ex. A-20 ¶¶ 41-43.) See *Uniloc 2017*, 2020 WL 569868 at *15.

As previously noted, the PTO has specifically instructed patentees to use the CRM claim form. A holding that CRM claims are necessarily means-plus-function, as Cirba proposes, would contradict the PTO’s explicit guidance and imperil the validity of thousands of duly-issued CRM claims using that form. No court has so held. See *Apple Inc. v. Motorola, Inc.*, 757 F.3d 1286, 1298-99 (Fed. Cir. 2014) (“Requiring traditional physical structure in software limitations lacking the term means would result in all of these limitations being construed as means-plus-function limitations and subsequently being found indefinite.”).

(ii) Even if the CRM Term Were Subject to § 112, ¶ 6, the ‘842 Patent Discloses Sufficient Structure

Even if the CRM term were subject to § 112, ¶ 6, the patent discloses sufficient structure. Claims 1 and 12 recite essentially the same method steps, differing only by claim 1’s recitation of the well-known DRS module. A POSITA understands that the same algorithm and related descriptions that support the distributed resource scheduling module term, described above, also support the CRM term. (’842 patent, 10:56-65 (“One or more embodiments of the present invention may be implemented as . . . one or more computer program modules embodied in one or more computer readable media.”); Ex. A-20 ¶¶ 44-47.)

d. Densify’s Sur-Reply Position

The CRM Instructions are subject to 35 U.S.C. §112, ¶ 6. The claims do not provide any information as to how the computer-readable medium is supposed to be programmed—they only describe the functions to be performed, not the algorithms by which they are performed. *Aristocrat Techs. Australia Pty Ltd. v. Int’l Game Tech.*, 521 F.3d 1328, 1334 (Fed. Cir. 2008) (“But that language simply describes the function to be performed, not the algorithm by which it

is performed.”); *see also Finisar Corp. v. DirecTV Grp., Inc.*, 523 F.3d 1323, 1340 (Fed. Cir. 2008) (“this passage provides ‘nothing more than a restatement of the function, as recited in the claim.’”). For example, a POSA would understand that the “determining” and “calculating” limitations require that certain things must be “determined,” but because the claim says nothing about *how* to make those determinations, the POSA is given no information as to how to go about “determining an impact” or “calculating a combined impact.” Ex. B-2 ¶¶ 45-52.

The Specification Fails to Disclose Sufficient Structure. The ’842 patent does not disclose sufficient structure. VMware contends claims 1 and 12 are essentially the same method steps, differing only in that claim 1 includes the DRS module and claim 12 the CRM instructions. VMware contends claim 1 sets forth particular computer-implemented instructions that must be carried out and how those instructions interact, but these instructions merely describe the functions to be performed – they say nothing about *how* the functions get carried out. Ex. B-2 ¶¶ 47-48; *see also Aristocrat Techs.*, 521 F.3d at 1334 (“But that language simply describes the function to be performed, not the algorithm by which it is performed.”) There is no disclosure of an algorithm in the claims, specification, or figures of the ’842 patent. Ex. B-2 ¶¶ 47-48.

3. “a virtual management computer configured to” (’842 Patent: claim 17)

VMware’s Proposed Construction	Densify’s Proposed Construction
Not a means-plus-function term subject to §112, ¶ 6 and not indefinite. As preambles, these terms are not limiting, merely identify well-known claim forms (non-transitory computer-readable storage media and computer systems), and should not be construed.	This is a means-plus-function term subject to § 112, ¶ 6. The term is indefinite for lack of sufficiently definite corresponding structure.

a. VMware’s Opening Position

VMware construed this term together with the “non-transient computer-readable storage medium” (claim 12) for its Opening Claim Construction Brief, *supra* Section III.D.2.a.

Following Cirba’s separate construction of these terms in its Answering Claim Construction Brief, VMware also construed these terms separately.

b. Densify’s Answering Position²⁹

The issues with this term mirror that of the term “non-transitory computer-readable storage medium comprising instructions...” and the outcome should be the same: the term is indefinite.

Claim 17 claims “one or more host computers” and “a virtual management computer, configured to” perform a series of steps identical to those recited in claim 12. The “virtual management computer” is the only structure identified for carrying out each and every step. But because “virtual management computer” is not a term of art, has no readily ascertainable meaning, and the claim provides no other structure for carrying out the steps, the term is a “black box” of the type criticized in *Williamson* and by this Court in *Arendi*, and Section 112 ¶ 6 applies. As with “non-transitory computer-readable storage medium comprising instructions...” because the specification provides inadequate disclosures (in fact, it provides no disclosures at all) as to the algorithm or algorithms for carrying out the steps, including at least the “calculating a combined impact” step, the term is indefinite. *Arendi*, 2019 WL 3891150 at *13; Ex. B-2 at ¶¶ 53-60.

²⁹ Densify previously asserted that “A computer system, comprising” was limiting; Densify withdraws that. Densify also previously included “one or more host computers” as part of its Section 112 ¶ 6 means plus function analysis, but withdraws that as well.

c. VMware’s Reply Position

Cirba argues the “virtual management computer configured to” term should be treated identically to the CRM term—*i.e.*, as limiting and subject to § 112, ¶ 6. (*Supra* § III.D.3.b at page 156.) But Cirba again provides no expert support for its position that the term lacks structure, and Dr. Madisetti again “assumed” otherwise. (Ex. B-2 ¶ 53 (Dr. Madisetti “asked by counsel for Densify to assume that a ‘virtual management computer, configured to’ is a means-plus-function term subject to interpretation under § 112, ¶ 6”).) A POSITA understands this claim to be structural for the same reason as claim 12 above: it specifies a general purpose computer configured to carry out a specific set of computer-implemented steps detailed by the claim and further described by the specification. (’842 patent, 3:15-5:61, claim 17, Figs. 1B-2B; Ex. A-20 ¶¶ 48-52.) And as for claims 1 and 12, even if this term were subject to § 112, ¶ 6, the ’842 patent discloses an algorithm and sufficient structure. (Ex. A-20 ¶¶ 53-54.)

d. Densify’s Sur-Reply Position

As discussed above with claim 12, “a virtual management computer, configured to” is subject to §112, ¶ 6 and tries to execute on a general-purpose computer configured to execute a specific set of computer-implemented steps detailed by the claims and the specification. The claimed functions are not algorithms, and the specification discloses none. *Aristocrat Techs.*, 521 F.3d at 1334 (“But that language simply describes the function to be performed, not the algorithm by which it is performed.”).

4. “a combined impact on current and future workload” (’842 Patent: claims 1, 12 and 17)³⁰

VMware’s Proposed Construction	Densify’s Proposed Construction
--------------------------------	---------------------------------

³⁰ Cirba originally proposed, and subsequently withdrew, the terms “an impact on current workload”/ “an impact on future workload” in its Answering Brief.

Not indefinite. These terms should not be construed as they take their plain and ordinary meaning.	Indefinite for failure to point out with particularity and distinctly claim the subject matter such that one of ordinary skill in the art would be reasonably apprised of the bounds of the asserted claims.
--	--

a. VMware's Opening Position

These phrases use ordinary English words that require no construction — “impact,” “current,” “future,” “combined,” and “workload.” The phrases are not indefinite.

The claim structure indicates the phrases take their plain and ordinary meaning. They appear in the second, third, and fourth recited steps of claims 1, 12, and 17. These steps relate to:

- determining the “impact” (i.e., effect) of performing the recommended change received in the first claim step (such as migrating a VM from one host to another) on, respectively, current and future workload in the virtual computing environment; and
- calculating a combined effect on current and future workload as recited in the claims.

(See '842 patent, 11:64-12:11; Ex. A-12 (*Am. Heritage Dictionary of the English Language* (4th ed. 2001)) at 425 (defining “impact” as “effect or impression of one thing on another”).) A “workload” is well understood in the art as “the amount of work which a person or computer has to do.” (Ex. A-4 (*Dictionary of Computing* (6th ed. 2010)) at 362.) Nothing in these steps or the words they use is ambiguous or requires construction.

The specification discloses an example confirming that the plain and ordinary meanings of these phrases apply. Figure 5 describes an exemplary “cost-benefit analysis to determine whether recommendations suggested by the VM management center should be performed.” ('842 patent, 9:27-30.) There, the DRS module “receives a recommended change,” “determines an impact on the current workload if that change is performed,” “determines an impact on future workload if the change is performed,” and “calculates a combined impact on current and future

workload.” (*Id.*, 9:36-52.) The patent also describes exemplary “impacts” including “improvement in the cluster imbalance metric” and “improved performance” resulting from performing a recommended change. (*Id.*, 9:15-17 (“In considering load-balancing moves, each move is scored based on its improvement in the cluster imbalance metric and on its risk-adjusted costs and benefits.”); 10:17-19 (“In addition, a cost-benefit analysis can be performed to determine whether the recommended change would likely result in improved performance.”).)

In Figure 5, the patent provides an example of how various “impacts” may be employed:

At step **510**, the DRS module determines whether the combined impact on current and future workload is above or below a threshold. In some embodiments, moves that hurt both current and future cluster imbalance are rejected. Moves that improve both current and future cluster imbalance are scored as the sum of current and future improvement, weighted by confidence in the future prediction of VMs on the source and target hosts. Moves that improve current, but hurt future cluster imbalance, could be scored as the improvement in current imbalance, with the idea that cost-benefit analysis will likely reject this move.

(*Id.*, 9:51-62; *see also id.*, 12:12-13 (claim 1).) This disclosure supports the plain and unambiguous meanings of the disputed phrases. As they use plain English words that a POSITA would understand readily, the Court need not construe them.

b. Densify’s Answering Position³¹

The patent’s disclosures regarding calculating “a combined impact on current and future workloads” do not provide a POSA with the requisite certainty as to what the term means. The only discussion of calculating combined impacts relates to step 508 of Fig. 5:

At step **508**, the DRS module calculates a combined impact on current and future workload. At step **510**, the DRS module determines whether the combined impact on current and future workload is above or below a threshold. *In some embodiments, moves that hurt both current and future cluster imbalance are*

³¹ Densify withdraws its initial construction of “indefiniteness” for the terms “an impact on current workload” and “an impact on future workload.”

rejected. Moves that improve both current and future cluster imbalance are scored as the sum of current and future improvement, weighted by confidence in the future prediction of VMs on the source and target hosts. Moves that improve current, but hurt future cluster imbalance, could be scored as the improvement in current imbalance, with the idea that cost-benefit analysis will likely reject this move. Moves that hurt current but improve future cluster imbalance are challenging to analyze; these moves need to be considered in light of how much benefit may be derived in the future (along with the confidence in the future) versus how much impact there is on current cluster imbalance.

'842 patent at 9:53-67 (emphasis added). In the scenario where current workload is improved but future workload is harmed, the patent teaches that the impacts should *not* be combined, *i.e.*, the future impact is not scored. *Id.* And in the scenario where both current and future workloads are harmed, the specification provides only vague information as to whether their impacts should be combined, and if so, how. *Id.* Given these opaque and inconsistent disclosures (combine? don't combine? how should they be determined in order to combine?), a POSA would not have been able to reasonably ascertain whether or how to "combine" impacts on current and future workloads in cases where one or both workloads are negatively impacted. Ex. B-2 at ¶¶ 61-63. A POSA would not be reasonably apprised of the bounds of the asserted claims 1, 12, and 17. *Infinity Computer Prods., Inc. v. Oki Data Americas, Inc.*, 987 F.3d 1053, 1059 (Fed. Cir. 2021); Ex. B-2 at ¶¶ 61-63.

c. VMware's Reply Position

A POSITA understands the meaning of the phrase "combined impact on current and future workload" with reasonable certainty. (Ex. A-20 ¶¶ 55-56.) The phrase is written in plain English and takes its plain meaning, consistent with the claim language and specification.

The claim structure reinforces the plain meaning of "combined impact," explaining precisely what is used to calculate the combined impact—the "determined impact on current

workload from performing a recommended change” and the “determined impact on future workload from performing a recommend change”—and requiring the result to be an “impact” comparable to a threshold. (’842 patent, claims 1, 12, 17; Ex. A-20 ¶ 56.) The specification also supports the phrase’s plain meaning by providing exemplary ways to calculate “combined” impacts. (*Supra* § III.D.4.a at page 158-59; ’842 patent, 9:15-17, 9:51-62, 10:17-19; Ex. A-20 ¶ 57.)

Cirba tries to confuse this plain meaning by manufacturing inconsistencies in the specification’s examples. (*Supra* § III.D.4.b at page 160 (“combine? don’t combine? how should they be determined in order to combine?”).) But these examples simply address different ways in which the impacts on current and future workloads might be considered. (’842 patent, 9:15-17, 9:51-62, 10:17-19; Ex. A-20 ¶¶ 57-58.) Cirba’s references to *unclaimed* examples, in which a *combined* impact is *not* calculated, are irrelevant. *Infinity Computer Products, Inc. v. Oki Data America, Inc.*, does not help Cirba, as the applicants here never took “materially inconsistent” positions about the meaning of “combined impact” during prosecution. 987 F.3d 1053, 1058-59 (Fed. Cir. 2021). The phrase “combined impact on current and future workloads” therefore either should not be construed or should be given its plain and ordinary meaning.

d. Densify’s Sur-Reply Position

VMware dismisses the confusion over what it means to calculate “a combined impact on current and future workload” by contending that the confusion centers only on “unclaimed” examples in the specification. *Supra* § III.D.4.c at page 158. Not so. The claim claims “calculating...a combined impact on current and future workload...if the recommended change is performed...” In context, this means that the claim requires making this calculation for every scenario where a recommended change is hypothetically performed – in other words, *if the recommended change is performed*, then the combined impact *must* be calculated. The patent

teaches several such scenarios where, *if the change is performed*, then certain things happen to the current and future workloads. In one scenario, the performed change would “improve both current and future cluster imbalance.” ’842 patent at 9:59-62. In another scenario, the performed change would “improve current, but hurt future cluster imbalance.” *Id.* at 9:59-62. And in a third scenario, the performed change would “hurt current but improve future cluster imbalance.” *Id.* at 9:62-67. The patent teaches how to calculate the combined impact only for first scenario – this is what VMware’s expert, Dr. Dinda points to (Ex. A-20 at ¶57). For the other two scenarios—which are also encompassed by the claim—the specification hints at a way to combine for one (“could be scored as the improvement in current imbalance, with the idea that cost-benefit analysis will likely reject this move”) (’842 patent at 9:59-62) but offers only vague and uncertain guidance for the other (“challenging to analyze”) (*id.* at 9:62-63). A POSA could not reasonably ascertain the scope of “a combined impact on current and future workload” based on the specification’s vague and incomplete disclosure. *Nautilus, Inc. v. Biosig Instruments, Inc.*, 572 U.S. 898, 901, 134 S. Ct. 2120, 2124, 189 L. Ed. 2d 37 (2014).

Dated: November 22, 2021

/s/ Robert M. Vrana

Anne Shea Gaza (No. 4093)
Robert M. Vrana (No. 5666)
Samantha G. Wilson (No. 5816)
YOUNG CONAWAY STARGATT &
TAYLOR, LLP
Rodney Square
1000 North King Street
Wilmington, DE 19801
(302) 571-6600
agaza@ycst.com
rvrana@ycst.com
swilson@ycst.com

Attorneys for VMware, Inc.

/s/ Kenneth L. Dorsney

Kenneth L. Dorsney (#3726)
kdorsney@morrisjames.com
MORRIS JAMES LLP
500 Delaware Avenue, Suite 1500
Wilmington, DE 19801
Telephone: (302) 888-6800

*Attorneys for Cirba Inc. (d/b/a DENSIFY)
and Cirba IP, Inc.*

28848029