UNITED STATES DISTRICT COURT
EASTERN DISTRICT OF TEXAS
TYLER DIVISION

REALTIME DATA LLC d/b/a IXO,

Plaintiff,

v.

RACKSPACE US, INC., NETAPP, INC., and SOLIDFIRE, LLC,

Defendants.

Case No. 6:16-cv-961

AMENDED COMPLAINT FOR PATENT INFRINGEMENT AGAINST RACKSPACE US, INC., NETAPP, INC. AND SOLIDFIRE, LLC

This is an action for patent infringement arising under the Patent Laws of the United States of America, 35 U.S.C. § 1 *et seq.* in which Plaintiff Realtime Data LLC d/b/a IXO ("Plaintiff," "Realtime," or "IXO") makes the following allegations against Defendants Rackspace US, Inc. ("Rackspace"), NetApp, Inc. ("NetApp"), and SolidFire, LLC ("SolidFire") (collectively, "Defendants"):

PARTIES

1. Realtime is a limited liability company organized under the laws of the State of New York. Realtime has places of business at 5851 Legacy Circle, Plano, Texas 75024, 1828 E.S.E. Loop 323, Tyler, Texas 75701, and 116 Croton Lake Road, Katonah, New York, 10536. Realtime has been registered to do business in Texas since May 2011. Since the 1990s, Realtime has researched and developed specific solutions for data compression, including, for example, those that increase the speeds at which data can be stored and accessed. As recognition of its innovations rooted in this technological field, Realtime holds 47 United States patents and has numerous pending patent applications. Realtime has licensed patents in this portfolio to many of the world's leading technology companies. The patents-in-suit relate to Realtime's development

of advanced systems and methods for fast and efficient data compression using numerous innovative compression techniques based on, for example, particular attributes of the data.

2. On information and belief, Defendant Rackspace US, Inc. ("Rackspace") is a Delaware corporation with its principal office at 1 Fanatical Pl, Windcrest, TX 78218. On information and belief, Rackspace US, Inc. can be served through its registered agent, Corporation Service Company d/b/a CSC-Lawyers Incorporating Service, 211 E. 7th Street, Suite 620, Austin, Texas 78701.

3. On information and belief, Defendant NetApp, Inc. ("NetApp") is a Delaware corporation with its principal office at 495 East Java Drive, Sunnyvale, California 94089. On information and belief, NetApp maintains places of business in Dallas, Texas and Houston, Texas. On information and belief, NetApp has been registered to do business in Texas since March 28, 1994, and can be served through its registered agent, Corporation Service Company d/b/a CSC-Lawyers Incorporating Service, 211 E. 7th Street, Suite 620, Austin, Texas 78701.

4. On information and belief, Defendant SolidFire, LLC ("SolidFire") is a Delaware limited liability company with its principal office at 1620 Pearl Street #200, Boulder, CO 80302. On information and belief, SolidFire can be served through its registered agent, Corporation Service Company, 2711 Centerville Rd Suite 400, Wilmington, DE 19808. On information and belief, SolidFire is 100% owned by Defendant NetApp as of February 2, 2016.

5. On information and belief, RackSpace and NetApp have entered into a commercial partnership "to deliver hosted storage and cloud services", including through the use of products that infringe Realtime's patents, as further described below. *See, e.g.*, http://blog.rackspace.com/netapp-announces-cloud-and-storage-partnership-with-rackspace/. On information and belief, pursuant to this commercial partnership, NetApp promotes the use of Rackspace cloud services on its own website. *See, e.g.*, http://solutionconnection.netapp.com/Rackspace ("Rackspace® (NYSE: RAX) is the global leader in hybrid cloud … Hundreds of thousands of customers look to Rackspace to deliver the best-fit infrastructure for their IT needs."). NetApp also identifies Rackspace as an "AltaVault

partner" cloud services provider in conjunction with the NetApp AltaVault product. *See, e.g.*, http://solutionconnection.netapp.com/altavault ("AltaVault seamlessly integrates with an organization's current backup software, and with all leading public and private cloud storage providers. AltaVault partners fall into three categories: … Public cloud providers … Rackspace."). As explained below, NetApp AltaVault infringes the Asserted Patents.

6. Furthermore, on information and belief, pursuant to this partnership, "Rackspace and NetApp … worked together to create a repeatable, approved, and tested process to integrate NetApp storage solutions into Rackspace Private Cloud footprints." *See, e.g.*, https://developer.rackspace.com/blog/rpc-and-netapp-hearts-block-storage/. As a result, Rackspace and NetApp offer a joint product, "Rackspace Private Cloud Powered by NetApp", which integrates NetApp's deduplication and compression features into Rackspace Private Cloud. *See, e.g.*, http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features."). RackSpace also offers dedicated Network Attached Storage ("NAS") powered by NetApp Data ONTAP. *See, e.g.*, https://www.rackspace.com/managed-hosting/data-storage/network-attached ("NAS offers you deduplication for improved performance and efficiency"). Accordingly, Rackspace and NetApp are properly joined in this action pursuant to 35 U.S.C. § 299.

7. Upon information and belief, NetApp acquired all of the outstanding ownership interests of SolidFire on February 2, 2016 and is the 100% owner of SolidFire. NetApp also markets the product "NetApp SolidFire" on the website solidfire.com. *See, e.g.*, http://www.solidfire.com/why-netapp-solidfire. As explained below, SolidFire products infringe

the Asserted Patents. Accordingly, NetApp and SolidFire are properly joined in this action pursuant to 35 U.S.C. § 299.

JURISDICTION AND VENUE

8. This action arises under the patent laws of the United States, Title 35 of the United States Code. This Court has original subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

9. This Court has personal jurisdiction over Defendant Rackspace in this action because Rackspace has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over Rackspace would not offend traditional notions of fair play and substantial justice. Defendant Rackspace, directly and through subsidiaries or intermediaries (including distributors, retailers, and others), has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents. Rackspace is registered to do business in the State of Texas, maintains its principal place of business in Windcrest, Texas, and has appointed Corporation Service Company d/b/a CSC-Lawyers Incorporating Service, 211 E. 7th Street, Suite 620, Austin, Texas 78701 as its agent for service of process.

10. This Court has personal jurisdiction over Defendant NetApp in this action because NetApp has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over NetApp would not offend traditional notions of fair play and substantial justice. Defendant NetApp, directly and through subsidiaries or intermediaries (including distributors, retailers, and others), has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents. NetApp has been registered to do business in the State of Texas since 1994, maintains places of business in Dallas, Texas and Houston, Texas, and has appointed Corporation Service Company d/b/a CSC-Lawyers Incorporating Service, 211 E. 7th Street, Suite 620, Austin, Texas 78701 as

its agent for service of process.

11. This Court has personal jurisdiction over Defendant SolidFire in this action because SolidFire has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over SolidFire would not offend traditional notions of fair play and substantial justice. Defendant SolidFire, directly and through subsidiaries or intermediaries (including distributors, retailers, and others), has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents.

12. Venue is proper in this district under 28 U.S.C. §§ 1391(b), 1391(c) and 1400(b). Rackspace and NetApp are registered to do business in Texas, and upon information and belief, Rackspace, NetApp, and SolidFire have transacted business in the Eastern District of Texas and have committed acts of direct and indirect infringement in the Eastern District of Texas. Rackspace also maintains its principal place of business in Texas, and upon information and belief, NetApp also maintains places of business in Texas.

COUNT I

INFRINGEMENT OF U.S. PATENT NO. 7,161,506

13. Plaintiff realleges and incorporates by reference paragraphs 1-12 above, as if fully set forth herein.

14. Plaintiff Realtime is the owner by assignment of United States Patent No. 7,161,506 (“the ‘506 patent”) entitled “Systems and methods for data compression such as content dependent data compression.” The ‘506 patent was duly and legally issued by the United States Patent and Trademark Office on January 9, 2007. A true and correct copy of the ‘506 patent, including its reexamination certificates, is included as Exhibit A.

NetApp ONTAP

15. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products and services that infringe the ‘506 patent, and

continues to do so. By way of illustrative example, these infringing products and services include, without limitation, NetApp's compression products and services, such as, *e.g.*, NetApp ONTAP v8 (since November 2015) and ONTAP 9 (collectively, "ONTAP") and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems (including but not limited to the NetApp AFF8000 series, including but not limited to the AFF8080 EX, AFF8060, and AFF8040) and hybrid disk FAS systems (including but not limited to NetApp FAS2500 series hybrid storage arrays, including but not limited to the FAS2554, FAS2552, and FAS2520, and NetApp FAS8000 Series hybrid storage systems, including but not limited to FAS8020, FAS8040, FAS8060, and FAS8080 EX systems), FlexPod® (including but not limited to FlexPod Datacenter, FlexPod Express, and FlexPod Select), ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the '506 patent (collectively, "Accused Instrumentality").

16. On information and belief, NetApp has directly infringed and continues to infringe the '506 patent, for example, through its own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the '506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

17. The Accused Instrumentality satisfies literally and/or under the doctrine of

equivalents the claim requirement "A computer implemented method for compressing data". This system minimizes the amount of data transmitted over a network and stored on a backup device. The Accused Instrumentality employs several data compression techniques to achieve this goal. *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

18. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially

the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

19. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-

Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

20. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing data compression with a single data compression encoder, if a data type of the data block is not identified". *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

21. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously

compressed and/or stored by the Accused Instrumentality were found not to literally meet the "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

22. On information and belief, NetApp also directly infringes and continues to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

23. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

24. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the ‘506 patent.

25. On information and belief, NetApp has had knowledge of the ‘506 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘506 patent and knew of its infringement, including by way of this lawsuit.

26. Upon information and belief, NetApp’s affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the ‘506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, NetApp touts the space-saving benefits of its inline data reduction technologies, including compression, deduplication, and compaction, to users of ONTAP. *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 (“All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case.”); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx (“Reduce your

data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.”). Thus, with knowledge of the ‘506 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘506 patent, knowing that such use constituted infringement of the ‘506 patent.

27. For similar reasons, NetApp also induces its customers to use the Accused Instrumentalities to infringe other claims of the ‘506 patent. NetApp specifically intended and was aware that these normal and customary activities would infringe the ‘506 patent. NetApp performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘506 patent, knowing that such use constitutes infringement of the ‘506 patent.

28. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the ‘506 patent pursuant to 35 U.S.C. § 271.

29. As a result of NetApp’s infringement of the ‘506 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

NetApp AltaVault / RackSpace Private Cloud Powered By NetApp

30. On information and belief, NetApp has made, used, offered for sale, sold and/or

imported into the United States NetApp products that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the NetApp AltaVault product and all versions and variations thereof since the issuance of the '506 patent. Upon information and belief, RackSpace has collaborated with NetApp to make AltaVault compatible with RackSpace's cloud services. *See*, *e.g.*, http://solutionconnection.netapp.com/altavault ("AltaVault seamlessly integrates with an organization's current backup software, and with all leading public and private cloud storage providers. AltaVault partners fall into three categories: … Public cloud providers … Rackspace.").

31. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace and NetApp combination products and/or services that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, "Rackspace Private Cloud Powered by NetApp", which integrates NetApp's deduplication and compression features into Rackspace Private Cloud. *See*, *e.g.*, http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**"). Upon information and belief, the deduplication / compression features enabled by NetApp in "Rackspace Private Cloud Powered by NetApp" function in a similar manner as the deduplication / compression features in NetApp AltaVault, and both AltaVault and "Rackspace Private Cloud Powered by NetApp", and all versions and variations thereof since the issuance of the '506 patent (collectively, "Accused Instrumentality") infringe the '506 patent.

32. On information and belief, NetApp has directly infringed and continues to infringe the ‘506 patent, for example, through its own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the ‘506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp’s customers.

33. On information and belief, Rackspace has directly infringed and continues to infringe the ‘506 patent, for example, through its own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the ‘506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, Rackspace uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused

Instrumentality to Rackspace's customers.

34. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "A computer implemented method for compressing data". This system minimizes the amount of data transmitted over a network and stored on a backup device. The Accused Instrumentality employs several data compression techniques to achieve this goal. *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf:

http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp;

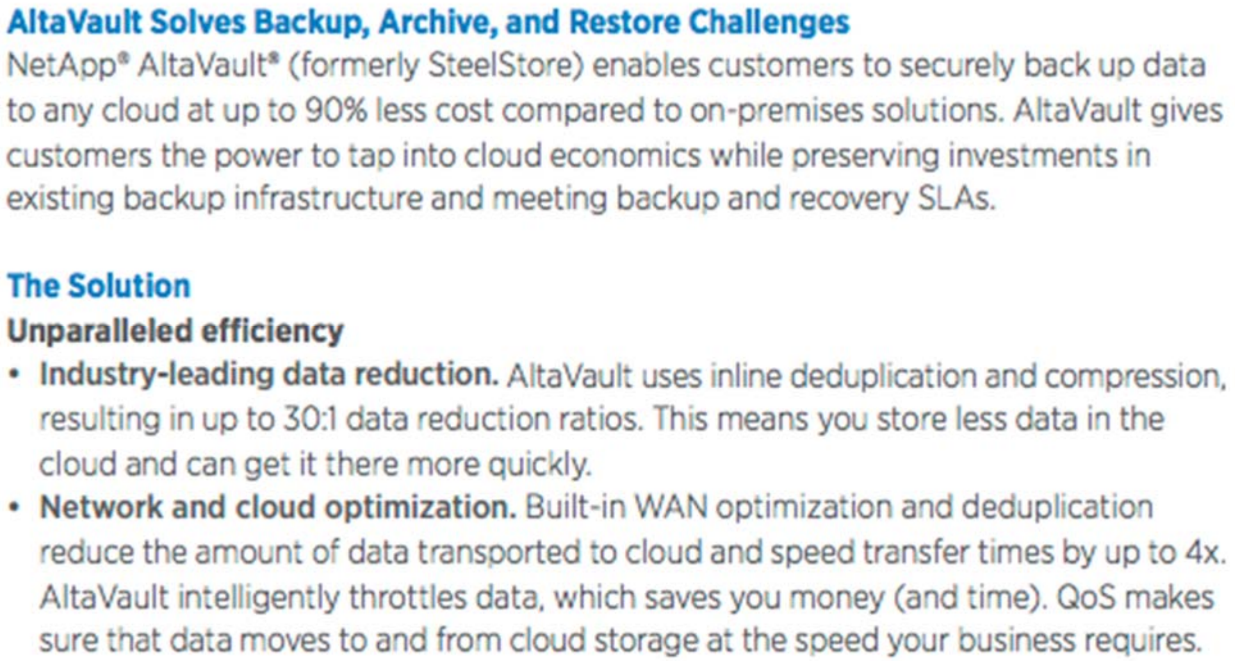

AltaVault Solves Backup, Archive, and Restore Challenges

NetApp® AltaVault® (formerly SteelStore) enables customers to securely back up data to any cloud at up to 90% less cost compared to on-premises solutions. AltaVault gives customers the power to tap into cloud economics while preserving investments in existing backup infrastructure and meeting backup and recovery SLAs.

The Solution

Unparalleled efficiency

- **Industry-leading data reduction.** AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly.
- **Network and cloud optimization.** Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x. AltaVault intelligently throttles data, which saves you money (and time). QoS makes sure that data moves to and from cloud storage at the speed your business requires.

http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

35. The Accused Instrumentality satisfies literally and/or under the doctrine of

equivalents the claim requirement "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

36. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform

de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

37. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing data compression with a single data compression encoder, if a data type of the data block is not identified". *See, e.g.*, https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no compression: Disables compression of any data written to the share. This is useful if you are copying over already-compressed data (for example: photos, videos, or proprietary formats such as medical data that might be compressed and encrypted already)."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

38. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of

the data type of the data within the data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp;
http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentI

D=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

39. On information and belief, NetApp and Rackspace also directly infringe and continue to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

40. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

41. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '506 patent.

42. On information and belief, NetApp and RackSpace have had knowledge of the '506 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp and Rackspace knew of the '506 patent and knew of their infringement, including by way of this lawsuit.

43. Upon information and belief, NetApp's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data

within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, NetApp emphasizes the benefits of the inline deduplication and compression features of NetApp AltaVault to its customers: "AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. … Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."). *See*, *e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf. NetApp also explains that its deduplication/compression approach is superior to others and results in faster restore times: "When you back up to AltaVault, it performs inline (real-time) deduplication of the backup data and replicates data into the cloud. AltaVault uses the local disk to store enough data for recovery of most recent backups. Such a mechanism provides LAN performance for the most likely restores. This deduplication process uses variable segment length inline deduplication plus compression, which is superior to other techniques such as fixed block. AltaVault deduplication level typically ranges between 10 and 30x. Deduplication performance depends on the incoming data type so turn off encryption and compression in the backup applications. Use the native encryption and deduplication in AltaVault to get higher data reduction rates than other typical software products. … AltaVault also optimizes restores from the cloud because it recalls only deduplicated data (which is not in the local cache) from the cloud. So if the customer is getting 10x deduplication, for example, and he or she needs to restore 10 TB of data, AltaVault needs only about 1 TB to restore. Over a 100-Mb line, this results in a time saving of days." See https://library.netapp.com/ecm/ecm_download_file/ECMP12434738. Thus, with knowledge of the '506 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '506 patent, knowing that such use constituted infringement of the '506 patent.

44. Upon information and belief, Rackspace's affirmative acts of making, using, and

selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, Rackspace instructs customers of the benefits of integrating NetApp storage solutions into Rackspace Private Cloud footprints: "Rackspace Private Cloud (RPC) powered by OpenStack … [can] connect your OpenStack cloud directly to a shared storage solution via Cinder integration drivers. This is where our friends at NetApp come into play. Rackspace and NetApp have formed a unique relationship to improve the Cinder shared storage capability within OpenStack. These two teams worked together to create a repeatable, approved, and tested process to integrate NetApp storage solutions into Rackspace Private Cloud footprints within a Rackspace datacenter or at the customer's datacenter." *See* https://developer.rackspace.com/blog/rpc-and-netapp-hearts-block-storage/. NetApp storage solutions used with Rackspace Private Cloud powered by NetApp include infringing deduplication / compression features. *See, e.g.*, http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.”). Rackspace specifically intended and was aware that these normal and customary activities would infringe the ‘506 patent. Thus, with knowledge of the ‘506 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘506 patent, knowing that such use constituted infringement of the ‘506 patent.

45. For similar reasons, NetApp and Rackspace also induce their customers to use the Accused Instrumentalities to infringe other claims of the ‘506 patent. NetApp and Rackspace specifically intended and were aware that these normal and customary activities would infringe the ‘506 patent. NetApp and Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp and Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp and Rackspace have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘506 patent, knowing that such use constitutes infringement of the ‘506 patent.

46. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp and Rackspace have injured Realtime and are liable to Realtime for infringement of the ‘506 patent pursuant to 35 U.S.C. § 271.

47. As a result of NetApp’s and Rackspace’s infringement of the ‘506 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp’s and Rackspace’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp and Rackspace, together with interest and costs as fixed by the Court.

<u>NetApp Mars OS</u>

48. On information and belief, NetApp has made, used, offered for sale, sold and/or

imported into the United States NetApp products that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the Mars OS software running on the FlashRay system and all versions and variations thereof since the issuance of the '506 patent ("Accused Instrumentality").

49. On information and belief, NetApp has directly infringed and continues to infringe the '506 patent, for example, through its own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the '506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

50. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "A computer implemented method for compressing data". This system minimizes the amount of data transmitted over a network and stored on a storage device. The Accused Instrumentality employs several data compression techniques to achieve this goal. *See, e.g.* https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to

implement data deduplication and compression more efficiently, thereby giving you more storage space savings.").

51. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

52. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing content dependent data compression with a content

dependent data compression encoder if a data type of the data block is identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

53. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing data compression with a single data compression encoder, if a data type of the data block is not identified". *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. … Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.").

54. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the analyzing of the data within the data block to

identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

55. On information and belief, NetApp also directly infringes and continues to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

56. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

57. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the ‘506 patent.

58. On information and belief, NetApp has had knowledge of the ‘506 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘506 patent and knew of its infringement, including by way of this lawsuit.

59. Upon information and belief, NetApp’s affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the ‘506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, NetApp instructs users of Mars OS about the benefits of its deduplication / compression features: “Industry-leading storage efficiency: Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.” *See* https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm.

Thus, with knowledge of the ‘506 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘506 patent, knowing that such use constituted infringement of the ‘506 patent.

60. For similar reasons, NetApp also induces its customers to use the Accused Instrumentalities to infringe other claims of the ‘506 patent. NetApp specifically intended and was aware that these normal and customary activities would infringe the ‘506 patent. NetApp performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘506 patent, knowing that such use constitutes infringement of the ‘506 patent.

61. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the ‘506 patent pursuant to 35 U.S.C. § 271.

62. As a result of NetApp’s infringement of the ‘506 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

Rackspace Cloud Backup

63. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace products that infringe the ‘506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Rackspace’s compression products and services, such as, e.g., Rackspace Cloud Backup and all

versions and variations thereof since the issuance of the ‘506 patent (“Accused Instrumentality”).

64. On information and belief, Rackspace has directly infringed and continues to infringe the ‘506 patent, for example, through its own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the ‘506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, Rackspace uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace’s customers.

65. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “A computer implemented method for compressing data”. This system minimizes the amount of data transmitted over a network and stored on a storage device. The Accused Instrumentality employs several data compression techniques to achieve this goal. *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 (“Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip.”); https://www.rackspace.com/en-us/cloud/backup (“Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup

includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

66. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality, for example, by identifying data matched or by comparing a hash value of a particular data block with a table of hash values for previously compressed and/or stored blocks, were found not to literally meet the "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available

data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

67. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified". Deduplication literally meets a correct construction of "compression" – "[Representing / represented / representation] of data with fewer bits". See Dkt. No. 362 in Case No. 6:15-cv-00463-RWS-JDL. Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example,

applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

68. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing data compression with a single data compression encoder, if a data type of the data block is not identified". *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup

includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

69. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality, for example, by identifying data matched or by comparing a hash value of a particular data block with a table of hash values for previously compressed and/or stored blocks, performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data

within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

70. On information and belief, Rackspace also directly infringes and continues to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

71. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

72. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '506 patent.

73. On information and belief, Rackspace has had knowledge of the '506 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, Rackspace knew of the '506 patent and knew of its infringement, including by way of this lawsuit.

74. Upon information and belief, Rackspace's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, Rackspace instructs users of Rackspace Cloud Backup about the benefits of its deduplication / compression features: "Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. Another benefit is that by using this method, you can retrieve previous versions of files, up to the limits specified by the customer-defined retention settings. De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a "delta" of the previous backup. These "deltas" are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required. To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip." See https://community.rackspace.com/products/f/25/t/7263. Rackspace also emphasizes the reduction in storage costs and backup time that the use of its deduplication/compression technology provides. See https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup,

additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."). Thus, with knowledge of the '506 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '506 patent, knowing that such use constituted infringement of the '506 patent.

75. For similar reasons, Rackspace also induces its customers to use the Accused Instrumentalities to infringe other claims of the '506 patent. Rackspace specifically intended and was aware that these normal and customary activities would infringe the '506 patent. Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Rackspace has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '506 patent, knowing that such use constitutes infringement of the '506 patent.

76. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Rackspace has injured Realtime and is liable to Realtime for infringement of the '506 patent pursuant to 35 U.S.C. § 271.

77. As a result of Rackspace's infringement of the '506 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Rackspace's infringement, but in no event less than a reasonable royalty for the use made of the invention by Rackspace, together with interest and costs as fixed by the Court.

SolidFire

78. On information and belief, NetApp and SolidFire (hereinafter,

"NetApp/SolidFire") have made, used, offered for sale, sold and/or imported into the United States NetApp/SolidFire products that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp/SolidFire's compression products and services, such as, *e.g.*, the SolidFire all-flash storage system and all versions and variations thereof since the issuance of the '506 patent ("Accused Instrumentality").

79. On information and belief, NetApp/SolidFire has directly infringed and continues to infringe the '506 patent, for example, through its own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the '506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, NetApp/SolidFire uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp/SolidFire's customers.

80. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "A computer implemented method for compressing data". This system minimizes the amount of data transmitted over a network and stored on a storage device. The Accused Instrumentality employs several data compression techniques to achieve this goal. *See, e.g.*, http://www.solidfire.com/platform/element-os/global-efficiencies ("With comprehensive data reduction from in-line deduplication, compression, thin provisioning, and space-efficient snapshots, SolidFire makes flash at scale an economic reality.").

81. The Accused Instrumentality satisfies literally and/or under the doctrine of

equivalents the claim requirement "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

82. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

83. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "performing data compression with a single data compression encoder, if a data type of the data block is not identified". *See, e.g.*,

http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire has architected its entire storage system accordingly, minimizing writes by compressing and deduplicating data before writing to its flash SSDs. When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. … If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value.").

84. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-

Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

85. On information and belief, NetApp/SolidFire also directly infringes and continues to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

86. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

87. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '506 patent.

88. On information and belief, NetApp/SolidFire has had knowledge of the '506 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp/SolidFire knew of the '506 patent and knew of its infringement, including by way of this lawsuit.

89. Upon information and belief, NetApp/SolidFire's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data

within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, NetApp/SolidFire instructs users of SolidFire about the benefits of its deduplication / compression features:

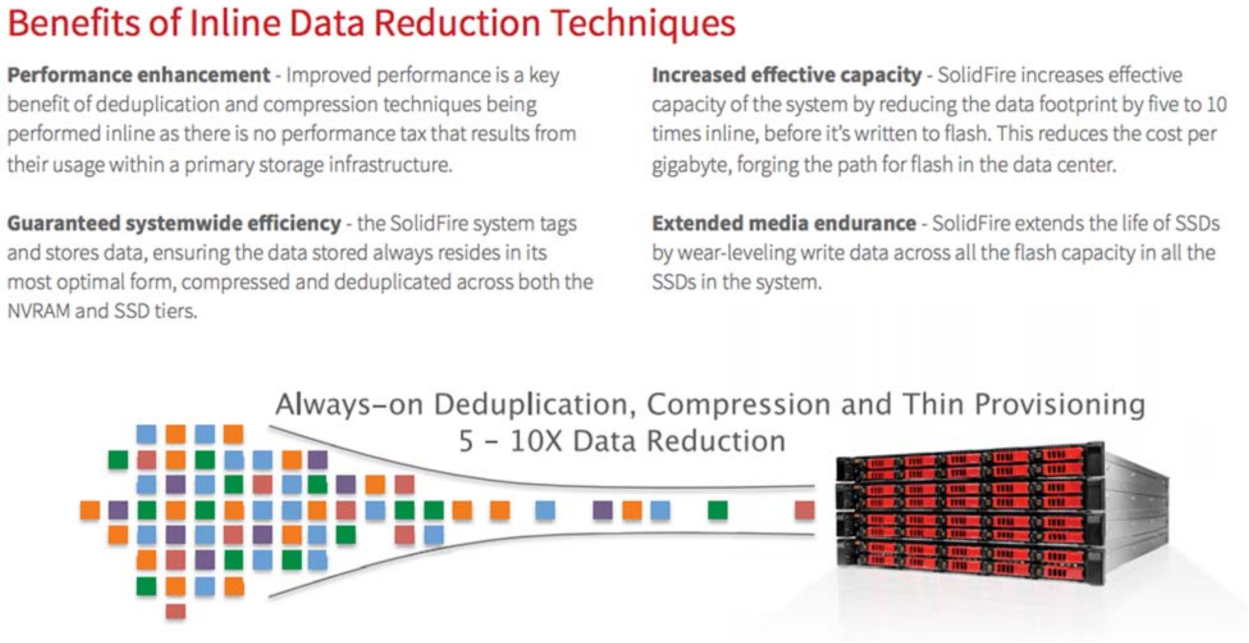

Benefits of Inline Data Reduction Techniques

Performance enhancement - Improved performance is a key benefit of deduplication and compression techniques being performed inline as there is no performance tax that results from their usage within a primary storage infrastructure.

Guaranteed systemwide efficiency - the SolidFire system tags and stores data, ensuring the data stored always resides in its most optimal form, compressed and deduplicated across both the NVRAM and SSD tiers.

Increased effective capacity - SolidFire increases effective capacity of the system by reducing the data footprint by five to 10 times inline, before it's written to flash. This reduces the cost per gigabyte, forging the path for flash in the data center.

Extended media endurance - SolidFire extends the life of SSDs by wear-leveling write data across all the flash capacity in all the SSDs in the system.

Always–on Deduplication, Compression and Thin Provisioning
5 – 10X Data Reduction

See http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf. Thus, with knowledge of the '506 patent gained from at least the filing and service of the original Complaint in this action, NetApp/SolidFire encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '506 patent, knowing that such use constituted infringement of the '506 patent.

90. For similar reasons, NetApp/SolidFire also induces its customers to use the Accused Instrumentalities to infringe other claims of the '506 patent. NetApp/SolidFire

specifically intended and was aware that these normal and customary activities would infringe the ‘506 patent. NetApp/SolidFire performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp/SolidFire engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp/SolidFire has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘506 patent, knowing that such use constitutes infringement of the ‘506 patent.

91. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp/SolidFire has injured Realtime and is liable to Realtime for infringement of the ‘506 patent pursuant to 35 U.S.C. § 271.

92. As a result of NetApp/SolidFire’s infringement of the ‘506 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp/SolidFire’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp/SolidFire, together with interest and costs as fixed by the Court.

COUNT II
INFRINGEMENT OF U.S. PATENT NO. 9,054,728

93. Plaintiff Realtime realleges and incorporates by reference paragraphs 1-92 above, as if fully set forth herein.

94. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,054,728 (“the ‘728 Patent”) entitled “Data compression systems and methods.” The ‘728 Patent was duly and legally issued by the United States Patent and Trademark Office on June 9, 2015. A true and correct copy of the ‘728 Patent is included as Exhibit B.

NetApp ONTAP

95. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products and services that infringe the ‘728 patent, and

continues to do so. By way of illustrative example, these infringing products and services include, without limitation, NetApp's compression products and services, such as, *e.g.*, NetApp ONTAP v8 (since November 2015) and ONTAP 9 (collectively, "ONTAP") and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems (including but not limited to the NetApp AFF8000 series, including but not limited to the AFF8080 EX, AFF8060, and AFF8040) and hybrid disk FAS systems (including but not limited to NetApp FAS2500 series hybrid storage arrays, including but not limited to the FAS2554, FAS2552, and FAS2520, and NetApp FAS8000 Series hybrid storage systems, including but not limited to FAS8020, FAS8040, FAS8060, and FAS8080 EX systems), FlexPod® (including but not limited to FlexPod Datacenter, FlexPod Express, and FlexPod Select), ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the '728 patent (collectively, "Accused Instrumentality").

96. On information and belief, NetApp has directly infringed and continues to infringe the '728 patent, for example, through its own use and testing of the Accused Instrumentality, which constitute systems for compressing data claimed by Claim 1 of the '728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, NetApp uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

97. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "A system for compressing data comprising; a processor; one or more content dependent data compression encoders". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "one or more content dependent data compression encoders" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

98. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a single data compression encoder." *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case.");

http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

99. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block". *See*, *e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

100. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of

the data are identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

101. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified". *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction

technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.”).

102. On information and belief, NetApp also directly infringes and continues to infringe other claims of the ‘728 patent, for similar reasons as explained above with respect to Claim 1 of the ‘728 patent.

103. On information and belief, all of the Accused Instrumentalities operate in substantially the same way.

104. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the ‘728 patent.

105. On information and belief, NetApp has had knowledge of the ‘728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘728 patent and knew of its infringement, including by way of this lawsuit.

106. Upon information and belief, NetApp’s affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the ‘728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, NetApp touts the space-saving benefits of its inline data reduction technologies, including compression, deduplication, and compaction, to users of

ONTAP. *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers."). Thus, with knowledge of the '728 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '728 patent, knowing that such use constituted infringement of the '728 patent.

107. For similar reasons, NetApp also induces its customers to use the Accused Instrumentalities to infringe other claims of the '728 patent. NetApp specifically intended and was aware that these normal and customary activities would infringe the '728 patent. NetApp performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '728 patent, knowing that such use constitutes infringement of the '728 patent.

108. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the '728 patent pursuant to 35 U.S.C. § 271.

109. As a result of NetApp's infringement of the '728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

NetApp AltaVault / RackSpace Private Cloud Powered By NetApp

110. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the '728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the NetApp AltaVault product and all versions and variations thereof since the issuance of the '728 patent. Upon information and belief, RackSpace has collaborated with NetApp to make AltaVault compatible with RackSpace's cloud services. *See, e.g.*, http://solutionconnection.netapp.com/altavault ("AltaVault seamlessly integrates with an organization's current backup software, and with all leading public and private cloud storage providers. AltaVault partners fall into three categories: … Public cloud providers … Rackspace.").

111. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace and NetApp combination products and/or services that infringe the '728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, "Rackspace Private Cloud Powered by NetApp", which integrates NetApp's deduplication and compression features into Rackspace Private Cloud. *See, e.g.*, http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's**

storage infrastructure enables the storage features.”). Upon information and belief, the deduplication / compression features enabled by NetApp in “Rackspace Private Cloud Powered by NetApp” function in a similar manner as the deduplication / compression features in NetApp AltaVault, and both AltaVault and “Rackspace Private Cloud Powered by NetApp”, and all versions and variations thereof since the issuance of the ‘728 patent (collectively, “Accused Instrumentality”) infringe the ‘728 patent.

112. On information and belief, NetApp has directly infringed and continues to infringe the ‘728 patent, for example, through its own use and testing of the Accused Instrumentality, which constitute systems for compressing data claimed by Claim 1 of the ‘728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, NetApp uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp’s customers.

113. On information and belief, Rackspace has directly infringed and continues to infringe the ‘728 patent, for example, through its own use and testing of the Accused Instrumentality, which constitute systems for compressing data claimed by Claim 1 of the ‘728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the

data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, Rackspace uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

114. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "A system for compressing data comprising; a processor; one or more content dependent data compression encoders". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "one or more content dependent data compression encoders" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform

de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

115. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a single data compression encoder." *See, e.g.*, https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no compression: Disables compression of any data written to the share. This is useful if you are copying over already-compressed data (for example: photos, videos, or proprietary formats such as medical data that might be compressed and encrypted already)."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

116. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes

analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block". *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

117. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store)

in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall. *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

118. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified". *See, e.g.*, https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no compression: Disables compression of any data written to the share. This is useful if you are copying over already-compressed data (for example: photos, videos, or proprietary formats such as medical data that might be compressed and encrypted already)."); http://solutionconnection.netapp.com/rackspace-

private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

119. On information and belief, NetApp and Rackspace also directly infringe and continue to infringe other claims of the '728 patent, for similar reasons as explained above with respect to Claim 1 of the '728 patent.

120. On information and belief, all of the Accused Instrumentalities operate in substantially the same way.

121. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '728 patent.

122. On information and belief, NetApp and Rackspace have had knowledge of the '728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, NetApp and Rackspace knew of the '728 patent and knew of their infringement, including by way of this lawsuit.

123. Upon information and belief, NetApp's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the '728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or

more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, NetApp emphasizes the benefits of the inline deduplication and compression features of NetApp AltaVault: "AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. … Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf. NetApp specifically intended and was aware that these normal and customary activities would infringe the '728 patent. Thus, with knowledge of the '728 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '728 patent, knowing that such use constituted infringement of the '728 patent.

124. Upon information and belief, Rackspace's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the '728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform

content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, Rackspace instructs customers of the benefits of integrating NetApp storage solutions into Rackspace Private Cloud footprints: "Rackspace Private Cloud (RPC) powered by OpenStack … [can] connect your OpenStack cloud directly to a shared storage solution via Cinder integration drivers. This is where our friends at NetApp come into play. Rackspace and NetApp have formed a unique relationship to improve the Cinder shared storage capability within OpenStack. These two teams worked together to create a repeatable, approved, and tested process to integrate NetApp storage solutions into Rackspace Private Cloud footprints within a Rackspace datacenter or at the customer's datacenter." *See* https://developer.rackspace.com/blog/rpc-and-netapp-hearts-block-storage/. NetApp storage solutions used with Rackspace Private Cloud powered by NetApp include infringing deduplication / compression features. *See, e.g.*, http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features."). Rackspace specifically intended and was aware that these normal and customary activities would infringe the '728 patent. Thus, with knowledge of the '728 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '728 patent, knowing that such use constituted infringement of the '728 patent.

125. For similar reasons, NetApp and Rackspace also induce their customers to use the

Accused Instrumentalities to infringe other claims of the ‘728 patent. NetApp and Rackspace specifically intended and were aware that these normal and customary activities would infringe the ‘728 patent. NetApp and Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp and Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp and Rackspace have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘728 patent, knowing that such use constitutes infringement of the ‘728 patent.

126. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp and Rackspace have injured Realtime and are liable to Realtime for infringement of the ‘728 patent pursuant to 35 U.S.C. § 271.

127. As a result of NetApp's and Rackspace's infringement of the ‘728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's and Rackspace's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp and Rackspace, together with interest and costs as fixed by the Court.

NetApp Mars OS

128. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the ‘728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the Mars OS software running on the FlashRay system and all versions and variations thereof since the issuance of the ‘728 patent (“Accused Instrumentality”).

129. On information and belief, NetApp has directly infringed and continues to infringe the ‘728 patent, for example, through its own use and testing of the Accused Instrumentality, which

constitute systems for compressing data claimed by Claim 1 of the ‘728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, NetApp uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp’s customers.

130. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “A system for compressing data comprising; a processor; one or more content dependent data compression encoders”. Even if the deduplication function in the Accused Instrumentality were found to not literally meet the “one or more content dependent data compression encoders” limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm (“Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.”).

131. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a single data compression encoder." *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. … Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.").

132. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block". *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

133. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication

performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

134. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified". *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. … Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.").

135. On information and belief, NetApp also directly infringes and continues to infringe other claims of the '728 patent, for similar reasons as explained above with respect to Claim 1 of the '728 patent.

136. On information and belief, all of the Accused Instrumentalities operate in substantially the same way.

137. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '728 patent.

138. On information and belief, NetApp has had knowledge of the '728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief,

NetApp knew of the ‘728 patent and knew of its infringement, including by way of this lawsuit.

139. Upon information and belief, NetApp’s affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the ‘728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, NetApp instructs users of Mars OS about the benefits of its deduplication / compression features: “Industry-leading storage efficiency: Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.” *See* https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm. Thus, with knowledge of the ‘728 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their

deduplication/compression functionality to infringe the ‘728 patent, knowing that such use constituted infringement of the ‘728 patent.

140. For similar reasons, NetApp also induces its customers to use the Accused Instrumentalities to infringe other claims of the ‘728 patent. NetApp specifically intended and was aware that these normal and customary activities would infringe the ‘728 patent. NetApp performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘728 patent, knowing that such use constitutes infringement of the ‘728 patent.

141. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the ‘728 patent pursuant to 35 U.S.C. § 271.

142. As a result of NetApp’s infringement of the ‘728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

Rackspace Cloud Backup

143. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace products that infringe the ‘728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Rackspace’s compression products and services, such as, e.g., Rackspace Cloud Backup and all versions and variations thereof since the issuance of the ‘728 patent (“Accused Instrumentality”).

144. On information and belief, Rackspace has directly infringed and continues to

infringe the ‘728 patent, for example, through its own use and testing of the Accused Instrumentality, which constitute systems for compressing data claimed by Claim 1 of the ‘728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, Rackspace uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace’s customers.

145. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “A system for compressing data comprising; a processor; one or more content dependent data compression encoders”. Deduplication literally meets a correct construction of “compression” – “[Representing / represented / representation] of data with fewer bits”. See Dkt. No. 362 in Case No. 6:15-cv-00463-RWS-JDL. Even if the deduplication function in the Accused Instrumentality were found to not literally meet the “one or more content dependent data compression encoders” limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 (“Rackspace Cloud Backup uses block-

level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

146. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a single data compression encoder." *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises

de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

147. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block". *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

148. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified". Deduplication literally meets a correct construction of "compression"

– "[Representing / represented / representation] of data with fewer bits". See Dkt. No. 362 in Case No. 6:15-cv-00463-RWS-JDL. Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

149. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified". *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-

level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip.”); https://www.rackspace.com/en-us/cloud/backup (“Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files.”); http://es.rack.ly/cloud/backup/how-it-works (“To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.”).

150. On information and belief, Rackspace also directly infringes and continues to infringe other claims of the ‘728 patent, for similar reasons as explained above with respect to Claim 1 of the ‘728 patent.

151. On information and belief, all of the Accused Instrumentalities operate in substantially the same way.

152. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the ‘728 patent.

153. On information and belief, Rackspace has had knowledge of the ‘728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, Rackspace knew of the ‘728 patent and knew of its infringement, including by way of this lawsuit.

154. Upon information and belief, Rackspace’s affirmative acts of making, using, and

selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the '728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, Rackspace instructs users of Rackspace Cloud Backup about the benefits of its deduplication / compression features: "Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. Another benefit is that by using this method, you can retrieve previous versions of files, up to the limits specified by the customer-defined retention settings. De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a "delta" of the previous backup. These "deltas" are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required. To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip." See https://community.rackspace.com/products/f/25/t/7263. Rackspace also emphasizes the reduction in storage costs and backup time that the use of its deduplication/compression technology provides. See https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files.”). Thus, with knowledge of the ‘728 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘728 patent, knowing that such use constituted infringement of the ‘728 patent.

155. For similar reasons, Rackspace also induces its customers to use the Accused Instrumentalities to infringe other claims of the ‘728 patent. Rackspace specifically intended and was aware that these normal and customary activities would infringe the ‘728 patent. Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Rackspace has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘728 patent, knowing that such use constitutes infringement of the ‘728 patent.

156. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, Rackspace has injured Realtime and is liable to Realtime for infringement of the ‘728 patent pursuant to 35 U.S.C. § 271.

157. As a result of Rackspace’s infringement of the ‘728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Rackspace’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Rackspace, together with interest and costs as fixed by the Court.

SolidFire

158. On information and belief, NetApp and SolidFire (hereinafter,

“NetApp/SolidFire”) have made, used, offered for sale, sold and/or imported into the United States NetApp/SolidFire products that infringe the ‘728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp/SolidFire’s compression products and services, such as, *e.g.*, the SolidFire all-flash storage system and all versions and variations thereof since the issuance of the ‘728 patent (“Accused Instrumentality”).

159. On information and belief, NetApp/SolidFire has directly infringed and continues to infringe the ‘728 patent, for example, through its own use and testing of the Accused Instrumentality, which constitute systems for compressing data claimed by Claim 1 of the ‘728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, NetApp/SolidFire uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp/SolidFire’s customers.

160. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “A system for compressing data comprising; a processor; one or more content dependent data compression encoders”. Even if the deduplication function in the Accused Instrumentality were found to not literally meet the “one or more content dependent data compression encoders” limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store)

in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

161. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a single data compression encoder." *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire has architected its entire storage system accordingly, minimizing writes by compressing and deduplicating data before writing to its flash SSDs. When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. … If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value.").

162. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes

analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block". *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

163. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the

unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.”).

164. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement “to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified”. *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf (“SolidFire has architected its entire storage system accordingly, minimizing writes by compressing and deduplicating data before writing to its flash SSDs. When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node’s NVRAM write cache. … If the compressed data block has unique data, it is stored in the system’s block pool. The block pool is organized only by the content hash value.”).

165. On information and belief, NetApp/SolidFire also directly infringes and continues to infringe other claims of the ‘728 patent, for similar reasons as explained above with respect to Claim 1 of the ‘728 patent.

166. On information and belief, all of the Accused Instrumentalities operate in substantially the same way.

167. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the ‘728 patent.

168. On information and belief, NetApp/SolidFire has had knowledge of the ‘728 patent

since at least the filing of the original Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘728 patent and knew of its infringement, including by way of this lawsuit.

169. Upon information and belief, NetApp/SolidFire’s affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the ‘728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, NetApp/SolidFire instructs users of SolidFire about the benefits of its deduplication / compression features:

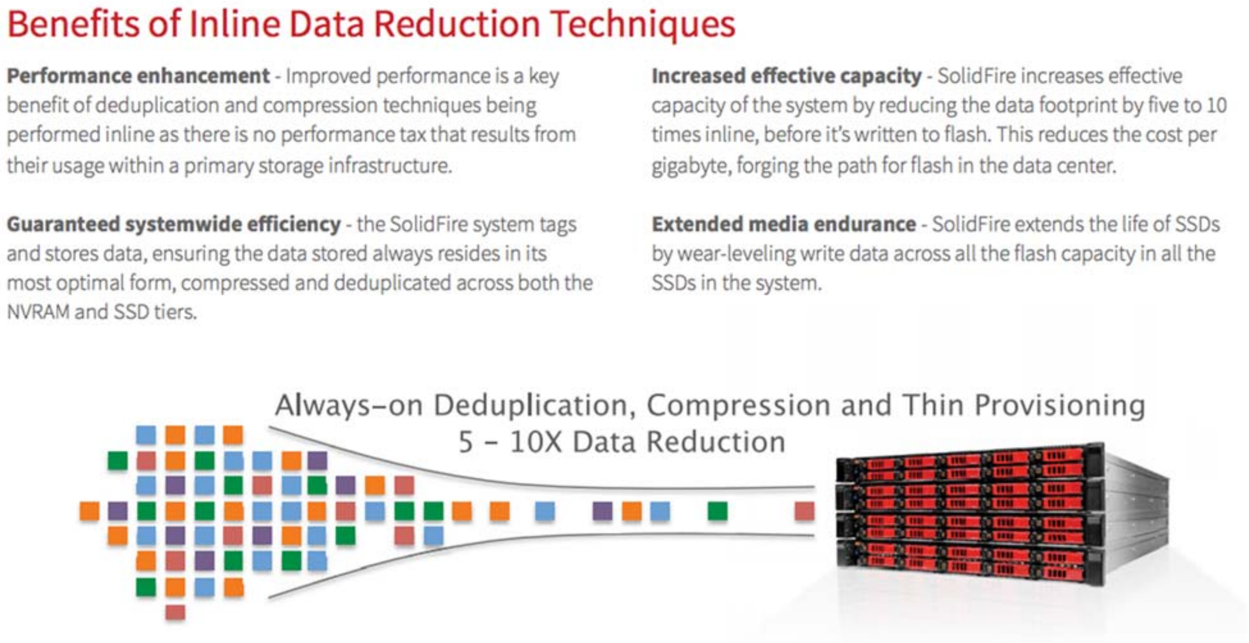

Benefits of Inline Data Reduction Techniques

Performance enhancement - Improved performance is a key benefit of deduplication and compression techniques being performed inline as there is no performance tax that results from their usage within a primary storage infrastructure.

Guaranteed systemwide efficiency - the SolidFire system tags and stores data, ensuring the data stored always resides in its most optimal form, compressed and deduplicated across both the NVRAM and SSD tiers.

Increased effective capacity - SolidFire increases effective capacity of the system by reducing the data footprint by five to 10 times inline, before it's written to flash. This reduces the cost per gigabyte, forging the path for flash in the data center.

Extended media endurance - SolidFire extends the life of SSDs by wear-leveling write data across all the flash capacity in all the SSDs in the system.

See http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf. Thus, with knowledge of the '728 patent gained from at least the filing and service of the original Complaint in this action, NetApp/SolidFire encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '728 patent, knowing that such use constituted infringement of the '728 patent.

170. For similar reasons, NetApp/SolidFire also induces its customers to use the Accused Instrumentalities to infringe other claims of the '728 patent. NetApp/SolidFire specifically intended and was aware that these normal and customary activities would infringe the '728 patent. NetApp/SolidFire performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp/SolidFire engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp/SolidFire has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and

customary way to infringe the ‘728 patent, knowing that such use constitutes infringement of the ‘728 patent.

171. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp/SolidFire has injured Realtime and is liable to Realtime for infringement of the ‘728 patent pursuant to 35 U.S.C. § 271.

172. As a result of NetApp/SolidFire’s infringement of the ‘728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp/SolidFire’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp/SolidFire, together with interest and costs as fixed by the Court.

COUNT III

INFRINGEMENT OF U.S. PATENT NO. 7,358,867

173. Plaintiff Realtime realleges and incorporates by reference paragraphs 1-172 above, as if fully set forth herein.

174. Plaintiff Realtime is the owner by assignment of United States Patent No. 7,358,867 (“the ‘867 Patent”) entitled “Content independent data compression method and system.” The ‘867 Patent was duly and legally issued by the United States Patent and Trademark Office on April 15, 2008. A true and correct copy of the ‘867 Patent is included as Exhibit C.

Rackspace Cloud Backup

175. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace products that infringe the ‘867 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Rackspace’s compression products and services, such as, e.g., Rackspace Cloud Backup and all versions and variations thereof since the issuance of the ‘867 patent (“Accused Instrumentality”).

176. On information and belief, Rackspace has directly infringed and continues to infringe the ‘867 patent, for example, through its own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 16 of the ‘867 patent, namely,

a method comprising: receiving a plurality of data blocks; determining whether or not to compress each one of said plurality of data blocks with a particular one or more of several encoders; if said determination is to compress with said particular one or more of said several encoders for a particular one of said plurality of data blocks; compressing said particular one of said plurality of data blocks with said particular one or more of said several encoders to provide a compressed data block; providing a data compression type descriptor representative of said particular one or more of said several encoders; outputting said data compression type descriptor and said compressed data block; if said determination is to not compress said particular one of said plurality of data blocks; providing a null data compression type descriptor representative of said determination not to compress; and outputting said null data compression type descriptor and said particular one of said plurality of data blocks. Upon information and belief, Rackspace uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

177. The Accused Instrumentality practices a method comprising: receiving a plurality of data blocks; determining whether or not to compress each one of said plurality of data blocks with a particular one or more of several encoders; if said determination is to compress with said particular one or more of said several encoders for a particular one of said plurality of data blocks; compressing said particular one of said plurality of data blocks with said particular one or more of said several encoders to provide a compressed data block; providing a data compression type descriptor representative of said particular one or more of said several encoders; outputting said data compression type descriptor and said compressed data block. *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup

includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

178. If said determination is to not compress said particular one of said plurality of data blocks, the Accused Instrumentality provides a null data compression type descriptor representative of said determination not to compress; and outputs said null data compression type descriptor and said particular one of said plurality of data blocks. *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size."). The Accused Instrumentality must utilize a null data compression type descriptor so that when retrieving the stored data, the Accused Instrumentality knows whether a given stored data block needs to be uncompressed or not in the process of data retrieval.

179. On information and belief, Rackspace also directly infringes and continues to infringe other claims of the '867 patent, for similar reasons as explained above with respect to Claim 16 of the '867 patent.

180. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

181. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the ‘867 patent.

182. On information and belief, Rackspace has had knowledge of the ‘867 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, Rackspace knew of the ‘867 patent and knew of its infringement, including by way of this lawsuit.

183. Upon information and belief, Rackspace’s affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the ‘867 patent by practicing a method comprising: receiving a plurality of data blocks; determining whether or not to compress each one of said plurality of data blocks with a particular one or more of several encoders; if said determination is to compress with said particular one or more of said several encoders for a particular one of said plurality of data blocks; compressing said particular one of said plurality of data blocks with said particular one or more of said several encoders to provide a compressed data block; providing a data compression type descriptor representative of said particular one or more of said several encoders; outputting said data compression type descriptor and said compressed data block; if said determination is to not compress said particular one of said plurality of data blocks; providing a null data compression type descriptor representative of said determination not to compress; and outputting said null data compression type descriptor and said particular one of said plurality of data blocks. For example, Rackspace instructs users of Rackspace Cloud Backup that, “Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip.” See https://community.rackspace.com/products/f/25/t/7263. Thus, with knowledge of the ‘867 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression

functionality to infringe the ‘867 patent, knowing that such use constituted infringement of the ‘867 patent.

184. Rackspace specifically intended and was aware that the normal and customary use of compression in the Accused Instrumentalities would infringe the ‘867 patent. Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘867 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Rackspace has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘867 patent, knowing that such use constitutes infringement of the ‘867 patent.

185. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, Rackspace has injured Realtime and is liable to Realtime for infringement of the ‘867 patent pursuant to 35 U.S.C. § 271.

186. As a result of Rackspace’s infringement of the ‘867 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Rackspace’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Rackspace, together with interest and costs as fixed by the Court.

COUNT IV

INFRINGEMENT OF U.S. PATENT NO. 7,378,992

187. Plaintiff realleges and incorporates by reference paragraphs 1-186 above, as if fully set forth herein.

188. Plaintiff Realtime is the owner by assignment of United States Patent No. 7,378,992 (“the ‘992 patent”) entitled “Content independent data compression method and system.” The ‘992 patent was duly and legally issued by the United States Patent and Trademark Office on May 27, 2008. A true and correct copy of the ‘992 patent, including its reexamination certificates, is

included as Exhibit D.

<u>NetApp ONTAP</u>

189. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products and services that infringe the '992 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, NetApp's compression products and services, such as, e.g., NetApp ONTAP v8 (since November 2015) and ONTAP 9 (collectively, "ONTAP") and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems (including but not limited to the NetApp AFF8000 series, including but not limited to the AFF8080 EX, AFF8060, and AFF8040) and hybrid disk FAS systems (including but not limited to NetApp FAS2500 series hybrid storage arrays, including but not limited to the FAS2554, FAS2552, and FAS2520, and NetApp FAS8000 Series hybrid storage systems, including but not limited to FAS8020, FAS8040, FAS8060, and FAS8080 EX systems), FlexPod® (including but not limited to FlexPod Datacenter, FlexPod Express, and FlexPod Select), ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the '992 patent (collectively, "Accused Instrumentality").

190. On information and belief, NetApp has directly infringed and continues to infringe at least claim 48 of the '992 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a

descriptor that is indicative of the data type of the data within the data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

191. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a computer implemented method comprising: receiving a data block". *See, e.g.,* https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

192. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "associating at least one encoder to each one of several data types." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "associating at least one encoder to each one of several data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple

available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

193. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "analyzing data within the data block to identify a first data type of the data within the data block". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "analyzing data within the data block to identify a first data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in

substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

194. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has

been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

195. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide

said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

196. On information and belief, NetApp also directly infringes and continues to infringe other claims of the '992 patent, for similar reasons as explained above with respect to Claim 48 of the '992 patent.

197. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

198. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '992 patent.

199. On information and belief, NetApp has had knowledge of the '992 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the '992 patent and knew of its infringement, including by way of this lawsuit.

200. NetApp's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the '992 patent by practicing compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, NetApp touts the space-saving benefits of its inline data reduction technologies, including compression, deduplication, and compaction, to users of ONTAP. *See*, *e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case.");

http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers."). Thus, with knowledge of the '992 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '992 patent, knowing that such use constituted infringement of the '992 patent.

201. For similar reasons, NetApp also induces its customers to use the Accused Instrumentalities to infringe other claims of the '992 patent. NetApp specifically intended and was aware that these normal and customary activities would infringe the '992 patent. NetApp performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '992 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '992 patent, knowing that such use constitutes infringement of the '992 patent.

202. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the '992 patent pursuant to 35 U.S.C. § 271.

203. As a result of NetApp's infringement of the '992 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

NetApp AltaVault / RackSpace Private Cloud Powered By NetApp

204. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the '992 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the NetApp AltaVault product, NetApp ONTAP v8 (since November 2015) and ONTAP 9 (collectively, "ONTAP") and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems and hybrid disk FAS systems, FlexPod®, ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the '992 patent. Upon information and belief, RackSpace has collaborated with NetApp to make AltaVault compatible with RackSpace's cloud services. *See*, *e.g.*, http://solutionconnection.netapp.com/altavault ("AltaVault seamlessly integrates with an organization's current backup software, and with all leading public and private cloud storage providers. AltaVault partners fall into three categories: … Public cloud providers … Rackspace.").

205. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace and NetApp combination products and/or services that infringe the '992 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, "Rackspace Private Cloud Powered by NetApp", which integrates NetApp's deduplication and compression features into Rackspace Private Cloud. *See*, *e.g.*, http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**"). Upon information and belief, the deduplication / compression features enabled by NetApp in "Rackspace Private Cloud Powered

by NetApp" function in a similar manner as the deduplication / compression features in NetApp AltaVault, and both AltaVault and "Rackspace Private Cloud Powered by NetApp", and all versions and variations thereof since the issuance of the '992 patent (collectively, "Accused Instrumentality") infringe the '992 patent.

206. On information and belief, NetApp has directly infringed and continues to infringe at least claim 48 of the '992 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

207. On information and belief, Rackspace has directly infringed and continues to infringe at least claim 48 of the '992 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block;

and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, Rackspace uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

208. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a computer implemented method comprising: receiving a data block". See, e.g., http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("AltaVault Solves Backup, Archive, and Restore Challenges: NetApp® AltaVault® (formerly SteelStore) enables customers to securely back up data to any cloud at up to 90% less cost compared to on-premises solutions. AltaVault gives customers the power to tap into cloud economics while preserving investments in existing backup infrastructure and meeting backup and recovery SLAs.");http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

209. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "associating at least one encoder to each one of several data types." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "associating at least one encoder to each one

of several data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.")

210. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "analyzing data within the data block to identify a first data type of the data within the data block". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "analyzing data within the data block to identify a first data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-

netapp;

http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.")

211. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused

Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.")

212. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide

said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication. no compression: Disables compression of any data written to the share. This is useful if you are copying over already-compressed data (for example: photos, videos, or proprietary formats such as medical data that might be compressed and encrypted already).");http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentI

D=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.")

213. On information and belief, NetApp and Rackspace also directly infringe and continue to infringe other claims of the '992 patent, for similar reasons as explained above with respect to Claim 48 of the '992 patent.

214. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

215. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '992 patent.

216. On information and belief, NetApp and Rackspace have had knowledge of the '992 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp and Rackspace knew of the '992 patent and knew of their infringement, including by way of this lawsuit.

217. NetApp's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the '992 patent by practicing compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data

block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, NetApp emphasizes the benefits of the inline deduplication and compression features of NetApp AltaVault to its customers: "AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. … Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf. Thus, with knowledge of the '992 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '992 patent, knowing that such use constituted infringement of the '992 patent.

218. Rackspace's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the '992 patent by practicing compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, Rackspace instructs customers of the benefits of integrating NetApp storage solutions into

Rackspace Private Cloud footprints: "Rackspace Private Cloud (RPC) powered by OpenStack … [can] connect your OpenStack cloud directly to a shared storage solution via Cinder integration drivers. This is where our friends at NetApp come into play. Rackspace and NetApp have formed a unique relationship to improve the Cinder shared storage capability within OpenStack. These two teams worked together to create a repeatable, approved, and tested process to integrate NetApp storage solutions into Rackspace Private Cloud footprints within a Rackspace datacenter or at the customer's datacenter." *See* https://developer.rackspace.com/blog/rpc-and-netapp-hearts-block-storage/. NetApp storage solutions used with Rackspace Private Cloud powered by NetApp include infringing deduplication / compression features. *See, e.g.*, http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features."). Thus, with knowledge of the '992 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '992 patent, knowing that such use constituted infringement of the '992 patent.

219. For similar reasons, NetApp and Rackspace also induce their customers to use the Accused Instrumentalities to infringe other claims of the '992 patent. NetApp and Rackspace specifically intended and were aware that these normal and customary activities would infringe the '992 patent. NetApp and Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '992 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp and Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp and Rackspace have

induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘992 patent, knowing that such use constitutes infringement of the ‘992 patent.

220. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp and Rackspace have injured Realtime and are liable to Realtime for infringement of the ‘992 patent pursuant to 35 U.S.C. § 271.

221. As a result of NetApp’s and Rackspace’s infringement of the ‘992 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp’s and Rackspace’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp and Rackspace, together with interest and costs as fixed by the Court.

NetApp Mars OS

222. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the ‘992 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp’s compression products and services, such as, *e.g.,* the Mars OS software running on the FlashRay system and all versions and variations thereof since the issuance of the ‘992 patent (“Accused Instrumentality”).

223. On information and belief, NetApp has directly infringed and continues to infringe at least claim 48 of the ‘992 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the ‘992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block

with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

224. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a computer implemented method comprising: receiving a data block". See, e.g., https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS is the software running on the FlashRay system, which stores data on an all-SSD (solid-state drive) shelf. … Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents.").

225. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "associating at least one encoder to each one of several data types." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "associating at least one encoder to each one of several data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused

Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

226. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "analyzing data within the data block to identify a first data type of the data within the data block". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "analyzing data within the data block to identify a first data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

227. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

228. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing, if said first data type is not the same as one of

said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.”).

229. On information and belief, NetApp also directly infringes and continues to infringe other claims of the ‘992 patent, for similar reasons as explained above with respect to Claim 48 of the ‘992 patent.

230. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

231. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the ‘992 patent.

232. On information and belief, NetApp has had knowledge of the ‘992 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘992 patent and knew of its infringement, including by way of this lawsuit.

233. NetApp’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the ‘992 patent by practicing compression methods claimed by the ‘992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, NetApp instructs users of Mars OS about the benefits of its deduplication / compression features: “Industry-

leading storage efficiency: Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs." *See* https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm. Thus, with knowledge of the '992 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '992 patent, knowing that such use constituted infringement of the '992 patent.

234. For similar reasons, NetApp also induces its customers to use the Accused Instrumentalities to infringe other claims of the '992 patent. NetApp specifically intended and was aware that these normal and customary activities would infringe the '992 patent. NetApp performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '992 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '992 patent, knowing that such use constitutes infringement of the '992 patent.

235. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the '992 patent pursuant to 35 U.S.C. § 271.

236. As a result of NetApp's infringement of the '992 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

Rackspace Cloud Backup

237. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace products that infringe the '992 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Rackspace's compression products and services, such as, e.g., Rackspace Cloud Backup and all versions and variations thereof since the issuance of the '992 patent ("Accused Instrumentality").

238. On information and belief, Rackspace has directly infringed and continues to infringe at least claim 48 of the '992 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, Rackspace uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

239. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a computer implemented method comprising: receiving a data

block". See, e.g., https://www.rackspace.com/en-us/cloud/backup ("Lost files and system failures can paralyze a project. That's why having a contingency plan is crucial. … File-level backups are transferred over our high-capacity network and written to three storage disks—all on separate nodes or locations that offer dual power supplies.").

240. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "associating at least one encoder to each one of several data types." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "associating at least one encoder to each one of several data types" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality, for example, by identifying data matched or by comparing a hash value of a particular data block with a table of hash values for previously compressed and/or stored blocks, performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, such as a hash value of a data block, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup

(“Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files.”); http://es.rack.ly/cloud/backup/how-it-works (“To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.”).

241. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “analyzing data within the data block to identify a first data type of the data within the data block”. Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the “analyzing data within the data block to identify a first data type of the data within the data block” limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality, for example, by identifying data matched or by comparing a hash value of a particular data block with a table of hash values for previously compressed and/or stored blocks, performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, such as a hash value of a data block, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the

optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

242. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block". Deduplication literally meets a correct construction of "compression" – "[Representing / represented / representation] of data with fewer bits". See Dkt. No. 362 in Case No. 6:15-cv-00463-RWS-JDL. Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining

whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality, for example, by identifying data matched or by comparing a hash value of a particular data block with a table of hash values for previously compressed and/or stored blocks, performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, such as a hash value of a data block, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

243. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing, if said first data type is not the same as one of

said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block." Deduplication literally meets a correct construction of "compression" – "[Representing / represented / representation] of data with fewer bits". See Dkt. No. 362 in Case No. 6:15-cv-00463-RWS-JDL. Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

244. On information and belief, Rackspace has had knowledge of the '992 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, Rackspace knew of the '992 patent and knew of its infringement, including by way of this lawsuit.

245. Rackspace's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the '992 patent by practicing compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the

data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, Rackspace instructs users of Rackspace Cloud Backup about the benefits of its deduplication / compression features: "Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. Another benefit is that by using this method, you can retrieve previous versions of files, up to the limits specified by the customer-defined retention settings. De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a "delta" of the previous backup. These "deltas" are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required. To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip." See https://community.rackspace.com/products/f/25/t/7263. Rackspace also emphasizes the reduction in storage costs and backup time that the use of its deduplication/compression technology provides. See https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."). Thus, with knowledge of the '992 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '992 patent, knowing that such use constituted infringement of the '992 patent.

246. For similar reasons, Rackspace also induces its customers to use the Accused Instrumentalities to infringe other claims of the '992 patent. Rackspace specifically intended and was aware that these normal and customary activities would infringe the '992 patent. Rackspace

performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '992 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Rackspace has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '992 patent, knowing that such use constitutes infringement of the '992 patent.

247. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Rackspace has injured Realtime and is liable to Realtime for infringement of the '992 patent pursuant to 35 U.S.C. § 271.

248. As a result of Rackspace's infringement of the '992 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Rackspace's infringement, but in no event less than a reasonable royalty for the use made of the invention by Rackspace, together with interest and costs as fixed by the Court.

SolidFire

249. On information and belief, NetApp and SolidFire (hereinafter, "NetApp/SolidFire") have made, used, offered for sale, sold and/or imported into the United States NetApp/SolidFire products that infringe the '992 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp/SolidFire's compression products and services, such as, *e.g.,* the SolidFire all-flash storage system and all versions and variations thereof since the issuance of the '992 patent ("Accused Instrumentality").

250. On information and belief, NetApp/SolidFire has directly infringed and continues to infringe at least claim 48 of the '992 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the '992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first

data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, NetApp/SolidFire uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp/SolidFire's customers.

251. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a computer implemented method comprising: receiving a data block". *See, e.g.*, http://www.solidfire.com/platform/element-os/global-efficiencies:

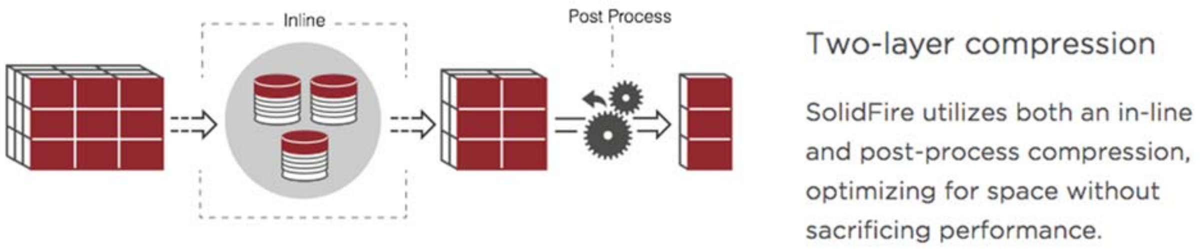

Two-layer compression

SolidFire utilizes both an in-line and post-process compression, optimizing for space without sacrificing performance.

252. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "associating at least one encoder to each one of several data types." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "associating at least one encoder to each one of several data types" limitation, this limitation is met under the doctrine of equivalents because it

is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

253. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "analyzing data within the data block to identify a first data type of the data within the data block". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "analyzing data within the data block to identify a first data type of the data within the data block" limitation, this

limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

254. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block". Even if the determination of whether particular data within a data block of an input data stream is

duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

255. The Accused Instrumentality satisfies literally and/or under the doctrine of

equivalents the claim requirement, “compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block.” Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the “compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block” limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf (“SolidFire’s Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. An acknowledgement is returned to the host when — and only when — the data has been safely stored in the NVRAM of multiple storage nodes. Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash. If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated.").

256. On information and belief, NetApp/SolidFire also directly infringes and continues to infringe other claims of the '992 patent, for similar reasons as explained above with respect to Claim 48 of the '992 patent.

257. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

258. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '992 patent.

259. On information and belief, NetApp/SolidFire has had knowledge of the '992 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp/SolidFire knew of the '992 patent and knew of its infringement, including by way of this lawsuit.

260. NetApp/SolidFire's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way

to infringe the ‘992 patent by practicing compression methods claimed by the ‘992 patent, including a computer implemented method comprising: receiving a data block; associating at least one encoder to each one of several data types; analyzing data within the data block to identify a first data type of the data within the data block; compressing if said first data type is the same as one of said several data types, said data block with said at least one encoder associated with said one of said several data types that is the same as said first data type to provide a compressed data block; and compressing, if said first data type is not the same as one of said several data types, said data block with a default encoder to provide said compressed data block, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, NetApp/SolidFire instructs users of SolidFire about the benefits of its deduplication / compression features:

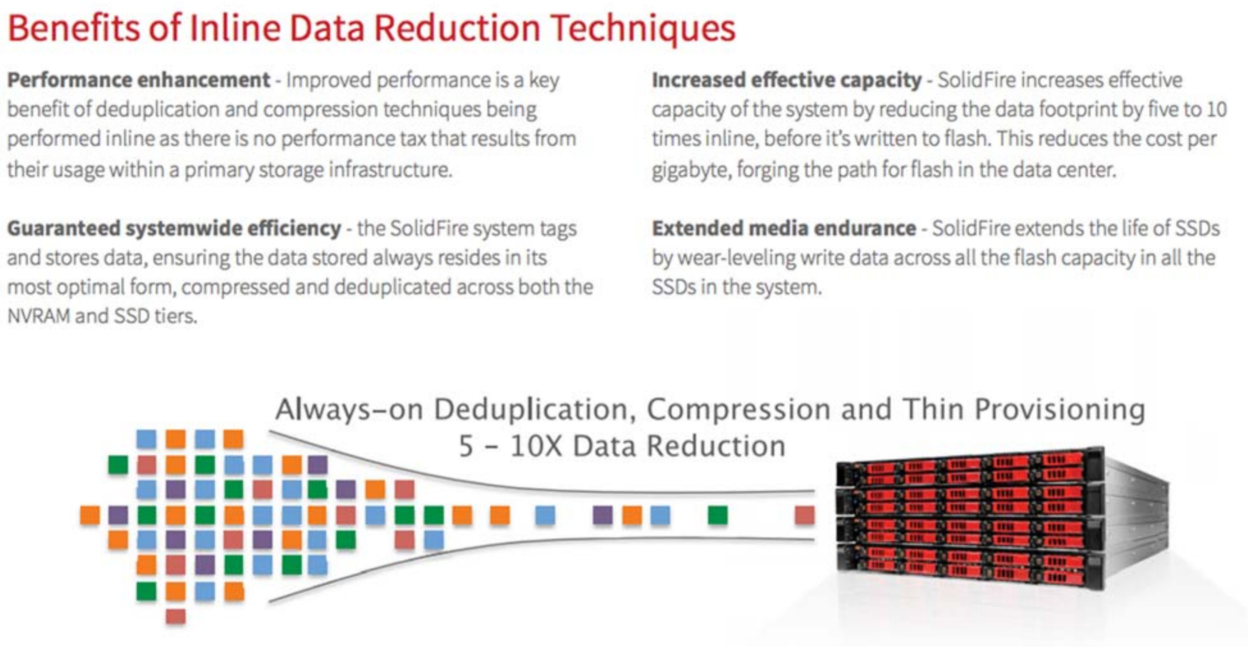

Benefits of Inline Data Reduction Techniques

Performance enhancement - Improved performance is a key benefit of deduplication and compression techniques being performed inline as there is no performance tax that results from their usage within a primary storage infrastructure.

Guaranteed systemwide efficiency - the SolidFire system tags and stores data, ensuring the data stored always resides in its most optimal form, compressed and deduplicated across both the NVRAM and SSD tiers.

Increased effective capacity - SolidFire increases effective capacity of the system by reducing the data footprint by five to 10 times inline, before it's written to flash. This reduces the cost per gigabyte, forging the path for flash in the data center.

Extended media endurance - SolidFire extends the life of SSDs by wear-leveling write data across all the flash capacity in all the SSDs in the system.

Always–on Deduplication, Compression and Thin Provisioning
5 – 10X Data Reduction

See also, e.g., http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf (“SolidFire’s Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the

4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node’s NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. An acknowledgement is returned to the host when — and only when — the data has been safely stored in the NVRAM of multiple storage nodes. Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash. If the compressed data block has unique data, it is stored in the system’s block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated.”). Thus, with knowledge of the ‘992 patent gained from at least the filing and service of the original Complaint in this action, NetApp/SolidFire encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘992 patent, knowing that such use constituted infringement of the ‘992 patent.

261. For similar reasons, NetApp/SolidFire also induces its customers to use the Accused Instrumentalities to infringe other claims of the ‘992 patent. NetApp/SolidFire specifically intended and was aware that these normal and customary activities would infringe the ‘992 patent. NetApp/SolidFire performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘992 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp/SolidFire engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp/SolidFire has induced and continues to induce

users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘992 patent, knowing that such use constitutes infringement of the ‘992 patent.

262. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp/SolidFire has injured Realtime and is liable to Realtime for infringement of the ‘992 patent pursuant to 35 U.S.C. § 271.

263. As a result of NetApp/SolidFire’s infringement of the ‘992 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp/SolidFire’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp/SolidFire, together with interest and costs as fixed by the Court.

COUNT V
INFRINGEMENT OF U.S. PATENT NO. 7,415,530

264. Plaintiff Realtime realleges and incorporates by reference paragraphs 1-263 above, as if fully set forth herein.

265. Plaintiff Realtime is the owner by assignment of United States Patent No. 7,415,530 (“the ‘530 Patent”) entitled “System and methods for accelerated data storage and retrieval.” The ‘530 Patent was duly and legally issued by the United States Patent and Trademark Office on August 19, 2008. A true and correct copy of the ‘530 Patent, including its reexamination certificate, is included as Exhibit E.

NetApp ONTAP

266. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products and services that infringe the ‘530 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, NetApp’s compression products and services, such as, *e.g.*, NetApp ONTAP v8 (since November 2015) and ONTAP 9 (collectively, “ONTAP”) and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems (including but not

limited to the NetApp AFF8000 series, including but not limited to the AFF8080 EX, AFF8060, and AFF8040) and hybrid disk FAS systems (including but not limited to NetApp FAS2500 series hybrid storage arrays, including but not limited to the FAS2554, FAS2552, and FAS2520, and NetApp FAS8000 Series hybrid storage systems, including but not limited to FAS8020, FAS8040, FAS8060, and FAS8080 EX systems), FlexPod® (including but not limited to FlexPod Datacenter, FlexPod Express, and FlexPod Select), ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the ‘530 patent (collectively, "Accused Instrumentality").

267. On information and belief, NetApp has directly infringed and continues to infringe at least Claim 1 of the ‘530 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

268. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "A system comprising: a memory device." *See, e.g.*,

https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

269. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream." *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

270. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "by compressing said first data block with a first compression

technique and said second data block with a second compression technique, said first and second compression techniques are different". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

271. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compressed data stream is stored on said memory device."

See, *e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

272. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form." *See*, *e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

273. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block." *See*, *e.g.*,

https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

274. On information and belief, NetApp also directly infringes and continues to infringe other claims of the '530 patent, for similar reasons as explained above with respect to Claim 1 of the '530 patent.

275. On information and belief, all of the Accused Instrumentalities constitute the claimed systems in substantially the same way.

276. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '530 patent.

277. On information and belief, NetApp has had knowledge of the '530 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the '530 patent and knew of its infringement, including by way of this lawsuit.

278. NetApp's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the '530 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data

accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block, thereby infringing the '530 patent. For example, NetApp touts the space-saving benefits of its inline data reduction technologies, including compression, deduplication, and compaction, to users of ONTAP. *See*, *e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers."). Thus, with knowledge of the '530 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '530 patent, knowing that such use constituted infringement of the '530 patent.

279. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp has injured Realtime and is liable to Realtime for infringement of

the '530 patent pursuant to 35 U.S.C. § 271.

280. As a result of NetApp's infringement of the '530 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

NetApp AltaVault / RackSpace Private Cloud Powered By NetApp

281. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the '530 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the NetApp AltaVault product, NetApp ONTAP v8 (since November 2015) and ONTAP 9 (collectively, "ONTAP") and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems and hybrid disk FAS systems, FlexPod®, ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the '530 patent. Upon information and belief, RackSpace has collaborated with NetApp to make AltaVault compatible with RackSpace's cloud services. *See*, *e.g.*, http://solutionconnection.netapp.com/altavault ("AltaVault seamlessly integrates with an organization's current backup software, and with all leading public and private cloud storage providers. AltaVault partners fall into three categories: … Public cloud providers … Rackspace.").

282. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace and NetApp combination products and/or services that infringe the '530 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, "Rackspace Private Cloud Powered by NetApp", which integrates NetApp's deduplication and compression features into Rackspace Private Cloud. *See*, *e.g.*, http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp

storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**"). Upon information and belief, the deduplication / compression features enabled by NetApp in "Rackspace Private Cloud Powered by NetApp" function in a similar manner as the deduplication / compression features in NetApp AltaVault, and both AltaVault and "Rackspace Private Cloud Powered by NetApp", and all versions and variations thereof since the issuance of the '530 patent (collectively, "Accused Instrumentality") infringe the '530 patent.

283. On information and belief, NetApp has directly infringed and continues to infringe at least Claim 1 of the '530 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

284. On information and belief, Rackspace has directly infringed and continues to infringe at least Claim 1 of the '530 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block. Upon information and belief, Rackspace uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

285. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "A system comprising: a memory device." *See, e.g.*, http://www.netapp.com/us/media/ds-3687.pdf ("AltaVault physical appliances are the industry's most scalable cloud-integrated storage appliances, with capacities ranging from 32TB up to 384TB of usable local cache. … AltaVault physical appliances are built on a scalable and efficient hardware platform that is optimized to reduce data footprints and rapidly stream data to the cloud."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering

maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

286. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream." *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf:

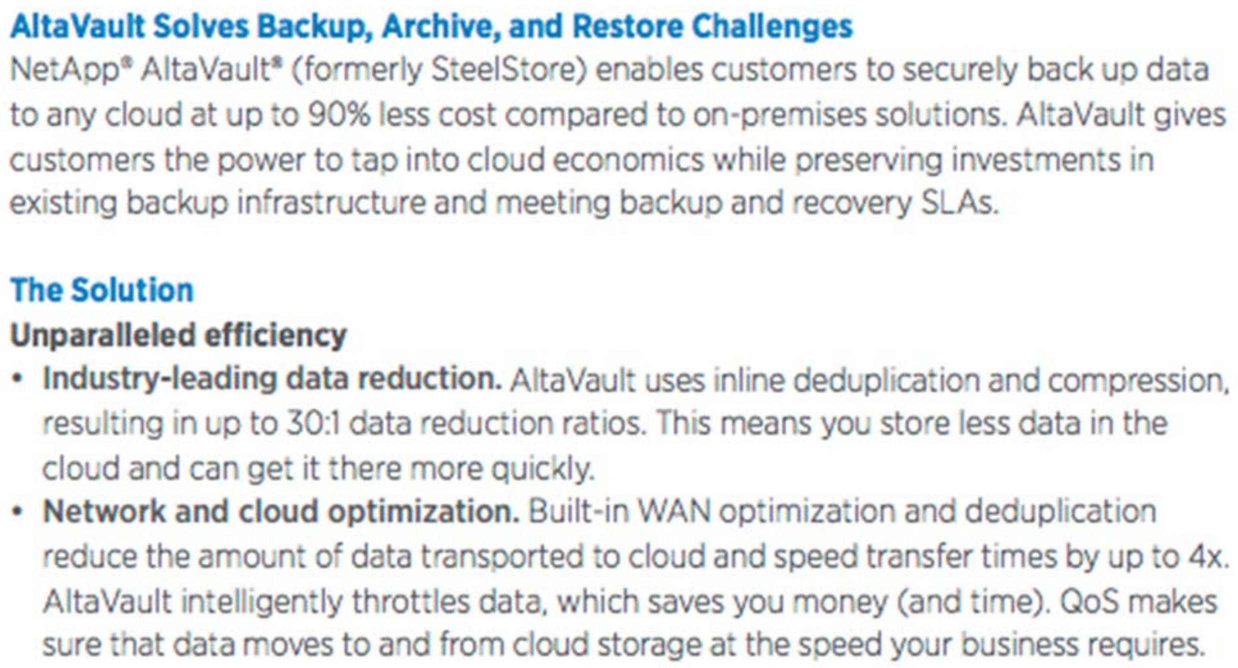

AltaVault Solves Backup, Archive, and Restore Challenges

NetApp® AltaVault® (formerly SteelStore) enables customers to securely back up data to any cloud at up to 90% less cost compared to on-premises solutions. AltaVault gives customers the power to tap into cloud economics while preserving investments in existing backup infrastructure and meeting backup and recovery SLAs.

The Solution

Unparalleled efficiency

- **Industry-leading data reduction.** AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly.
- **Network and cloud optimization.** Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x. AltaVault intelligently throttles data, which saves you money (and time). QoS makes sure that data moves to and from cloud storage at the speed your business requires.

http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

287. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp;

http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentI

D=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

288. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compressed data stream is stored on said memory device." *See*, *e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf:

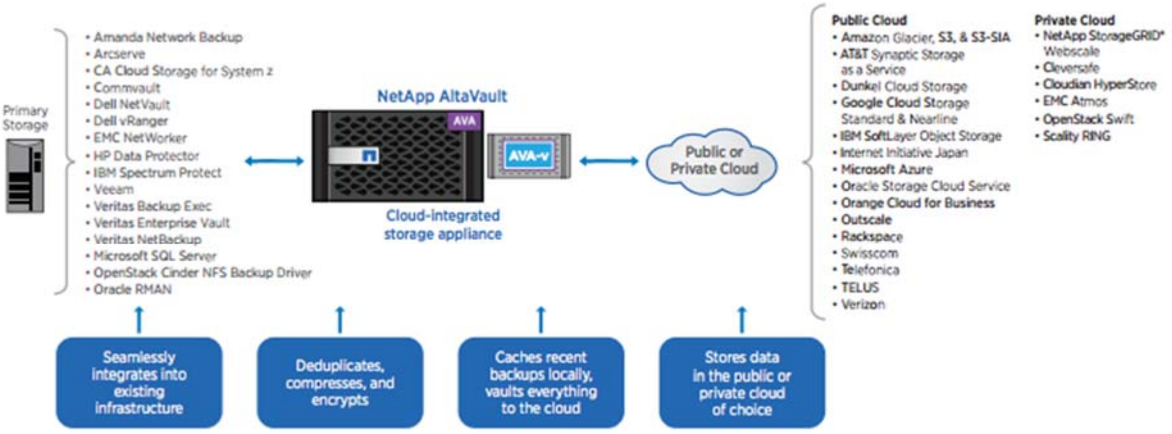
Primary Storage
• Amanda Network Backup
• Arcserve
• CA Cloud Storage for System z
• Commvault
• Dell NetVault
• Dell vRanger
• EMC NetWorker
• HP Data Protector
• IBM Spectrum Protect
• Veeam
• Veritas Backup Exec
• Veritas Enterprise Vault
• Veritas NetBackup
• Microsoft SQL Server
• OpenStack Cinder NFS Backup Driver
• Oracle RMAN
NetApp AltaVault
AVA
AVA-v
Cloud-integrated storage appliance
Public or Private Cloud
Public Cloud
• Amazon Glacier, S3, & S3-SIA
• AT&T Synaptic Storage as a Service
• Dunkel Cloud Storage
• Google Cloud Storage Standard & Nearline
• IBM SoftLayer Object Storage
• Internet Initiative Japan
• Microsoft Azure
• Oracle Storage Cloud Service
• Orange Cloud for Business
• Outscale
• Rackspace
• Swisscom
• Telefonica
• TELUS
• Verizon
Private Cloud
• NetApp StorageGRID* Webscale
• Cleversafe
• Cloudian HyperStore
• EMC Atmos
• OpenStack Swift
• Scality RING
Seamlessly integrates into existing infrastructure
Deduplicates, compresses, and encrypts
Caches recent backups locally, vaults everything to the cloud
Stores data in the public or private cloud of choice

Figure 1) NetApp AltaVault provides seamless integrations with existing applications and cloud service providers.

http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

289. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form." *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x. AltaVault intelligently throttles data, which saves you money (and time)."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

290. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block." *See, e.g.*, https://library.netapp.com/ecm/ecm_download_file/ECMP12434738 at 7 ("AltaVault also optimizes restores from the cloud because it recalls only deduplicated data (which is not in the local cache) from the cloud. So if the customer is getting 10x deduplication, for example, and he or she needs to restore 10 TB of data, AltaVault needs only about 1 TB to restore. Over a 100-Mb line, this results in a time saving of days. Data moves from the backup client to the backup server, to AltaVault, and then to the cloud. When you restore data, data moves from the cache in AltaVault, in which it is expanded to its original size to the backup server and to the backup client. If the data is not local, it moves from the cloud to AltaVault, to the backup server, and to the

backup client."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

291. On information and belief, NetApp and Rackspace also directly infringe and continue to infringe other claims of the '530 patent, for similar reasons as explained above with respect to Claim 1 of the '530 patent.

292. On information and belief, all of the Accused Instrumentalities constitute the claimed system in substantially the same way.

293. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '530 patent.

294. On information and belief, NetApp and RackSpace have had knowledge of the '530 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp and Rackspace knew of the '530 patent and knew of their infringement, including by way of this lawsuit.

295. NetApp's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the '530 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data

accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block, thereby infringing the '530 patent. For example, NetApp emphasizes the benefits of the inline deduplication and compression features of NetApp AltaVault to its customers: "AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. … Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf. NetApp also explains that its deduplication/compression approach is superior to others and results in faster restore times: "When you back up to AltaVault, it performs inline (real-time) deduplication of the backup data and replicates data into the cloud. AltaVault uses the local disk to store enough data for recovery of most recent backups. Such a mechanism provides LAN performance for the most likely restores. This deduplication process uses variable segment length inline deduplication plus compression, which is superior to other techniques such as fixed block. AltaVault deduplication level typically ranges between 10 and 30x. Deduplication performance depends on the incoming data type so turn off encryption and compression in the backup applications. Use the native encryption and deduplication in AltaVault to get higher data reduction rates than other typical software products. … AltaVault also optimizes restores from the cloud because it recalls only deduplicated data (which is not in the local cache) from the cloud. So if the customer is getting 10x deduplication, for example, and he or she needs to restore 10 TB of data, AltaVault needs only about 1 TB to restore. Over a 100-Mb line, this results in a time saving of days." See https://library.netapp.com/ecm/ecm_download_file/ECMP12434738. Thus, with knowledge of the '530 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '530 patent, knowing that such use constituted infringement of the '530 patent.

296. Rackspace's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the '530 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block, thereby infringing the '530 patent. For example, Rackspace instructs customers of the benefits of integrating NetApp storage solutions into Rackspace Private Cloud footprints: "Rackspace Private Cloud (RPC) powered by OpenStack … [can] connect your OpenStack cloud directly to a shared storage solution via Cinder integration drivers. This is where our friends at NetApp come into play. Rackspace and NetApp have formed a unique relationship to improve the Cinder shared storage capability within OpenStack. These two teams worked

together to create a repeatable, approved, and tested process to integrate NetApp storage solutions into Rackspace Private Cloud footprints within a Rackspace datacenter or at the customer's datacenter." *See* https://developer.rackspace.com/blog/rpc-and-netapp-hearts-block-storage/. NetApp storage solutions used with Rackspace Private Cloud powered by NetApp include infringing deduplication / compression features. *See, e.g.*, http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features."). Rackspace specifically intended and was aware that these normal and customary activities would infringe the '530 patent. Thus, with knowledge of the '530 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '530 patent, knowing that such use constituted infringement of the '530 patent.

297. For similar reasons, NetApp and Rackspace also induce their customers to use the Accused Instrumentalities to infringe other claims of the '530 patent. NetApp and Rackspace specifically intended and were aware that these normal and customary activities would infringe the '530 patent. NetApp and Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '530 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp and Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp and Rackspace have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '530 patent, knowing that

such use constitutes infringement of the ‘530 patent.

298. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp and Rackspace have injured Realtime and are liable to Realtime for infringement of the ‘530 patent pursuant to 35 U.S.C. § 271.

299. As a result of NetApp’s and Rackspace’s infringement of the ‘530 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp’s and Rackspace’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp and Rackspace, together with interest and costs as fixed by the Court.

NetApp Mars OS

300. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the ‘530 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp’s compression products and services, such as, *e.g.*, the Mars OS software running on the FlashRay system and all versions and variations thereof since the issuance of the ‘530 patent (“Accused Instrumentality”).

301. On information and belief, NetApp has directly infringed and continues to infringe at least Claim 1 of the ‘530 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on

said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

302. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "A system comprising: a memory device." *See, e.g.*, http://www.netapp.com/us/company/news/press-releases/news-rel-20140917-816753.aspx ("NetApp (NASDAQ: NTAP) today announced its first shipments of FlashRay™, the company's purpose-built all-flash storage array with the new NetApp® Mars™ operating system. FlashRay is designed from the ground up to improve the performance, efficiency, and manageability of all-flash storage architectures used in enterprise application environments."); https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS is the software running on the FlashRay system, which stores data on an all-SSD (solid-state drive) shelf.").

303. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream." *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only

a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.”).

304. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement “by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different”. Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the “compression technique” limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm (“Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only

a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.").

305. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compressed data stream is stored on said memory device." *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS is the software running on the FlashRay system, which stores data on an all-SSD (solid-state drive) shelf. … Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings.").

306. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form." *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS delivers high performance in the form of consistent and predictable low latency and high throughput. For small random access workloads, throughput is measured in I/O operations per second (IOPS); for sequential and large random workloads, it is measured as bandwidth (for example, GB/s).").

307. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block." *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data

deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.").

308. On information and belief, NetApp also directly infringes and continues to infringe other claims of the '530 patent, for similar reasons as explained above with respect to Claim 1 of the '530 patent.

309. On information and belief, all of the Accused Instrumentalities constitute the claimed systems in substantially the same way.

310. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '530 patent.

311. On information and belief, NetApp has had knowledge of the '530 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the '530 patent and knew of its infringement, including by way of this lawsuit.

312. NetApp's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the '530 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs

faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block, thereby infringing the ‘530 patent. For example, NetApp instructs users of Mars OS about the benefits of its deduplication / compression features: “Industry-leading storage efficiency: Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.” *See* https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm. Thus, with knowledge of the ‘530 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘530 patent, knowing that such use constituted infringement of the ‘530 patent.

313. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the ‘530 patent pursuant to 35 U.S.C. § 271.

314. As a result of NetApp’s infringement of the ‘530 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp’s infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

Rackspace Cloud Backup

315. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace products that infringe the ‘530 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Rackspace’s compression products and services, such as, e.g., Rackspace Cloud Backup and all versions and variations thereof since the issuance of the ‘530 patent (“Accused Instrumentality”).

316. On information and belief, Rackspace has directly infringed and continues to infringe at least Claim 1 of the ‘530 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block. Upon information and belief, Rackspace uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace’s customers.

317. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement “A system comprising: a memory device.” *See, e.g.*, https://www.rackspace.com/cloud/servers/features (“Cloud Servers reside in our world-class data centers, with ECC memory, and fully redundant networking and power all the way to the host. All virtual servers are backed by hardware RAID 10 storage.”).

318. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream." *See*, *e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a "delta" of the previous backup. These "deltas" are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required. To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

319. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "by compressing said first data block with a first compression

technique and said second data block with a second compression technique, said first and second compression techniques are different". Deduplication literally meets a correct construction of "compression" – "[Representing / represented / representation] of data with fewer bits". See Dkt. No. 362 in Case No. 6:15-cv-00463-RWS-JDL. Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a "delta" of the previous backup. These "deltas" are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required. To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

320. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compressed data stream is stored on said memory device." After deduplication and compression, the compressed data blocks are stored on the memory device of the Accused Instrumentality. See https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs.").

321. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form." *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a "delta" of the previous backup. These "deltas" are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required."); https://www.rackspace.com/en-gb/cloud/backup/features ("Cloud Backup includes file compression and block-level deduplication to reduce backup times, size, and storage costs. After the initial backup of a data set, block-level deduplication saves time and storage space by only

backing up files that have changed since the last backup.”).

322. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block.” The Accused Instrumentality must store a data descriptor that is used when retrieving a stored data block to distinguish between block requiring decompression and blocks that were not compressed before storage. *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 (“Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip.”); http://es.rack.ly/cloud/backup/how-it-works (“Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.”); https://community.rackspace.com/products/f/25/t/7263 (“Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead.”); https://support.rackspace.com/how-to/best-practices-for-cloud-backup/ (“Do not compress your data before it is backed up. Doing so defeats the backup deduplication, which is typically far more efficient than simple file compression. Deduplication stores only the updated data, and saves you storage space and money during the backup process.”); http://www.rackspace.co.uk/cloud/backup/how-it-works (“Built-in compression and de-duplication reduce the amount of data stored, while increasing the speed of subsequent backups and restores and reducing costs.”).

323. On information and belief, Rackspace also directly infringes and continues to infringe other claims of the ‘530 patent, for similar reasons as explained above with respect to Claim 1 of the ‘530 patent.

324. On information and belief, all of the Accused Instrumentalities constitute the

claimed systems in substantially the same way.

325. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the ‘530 patent.

326. On information and belief, Rackspace has had knowledge of the ‘530 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, Rackspace knew of the ‘530 patent and knew of its infringement, including by way of this lawsuit.

327. Rackspace’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the ‘530 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block, thereby infringing the ‘530 patent. For example, Rackspace instructs users of Rackspace Cloud Backup about the benefits of its deduplication / compression features: “Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. Another benefit is that by using this method, you can retrieve previous versions of files, up to the limits

specified by the customer-defined retention settings. De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a "delta" of the previous backup. These "deltas" are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required. To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip." See https://community.rackspace.com/products/f/25/t/7263. Rackspace also emphasizes the reduction in storage costs and backup time that the use of its deduplication/compression technology provides. See https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."). Thus, with knowledge of the '530 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '530 patent, knowing that such use constituted infringement of the '530 patent.

328. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Rackspace has injured Realtime and is liable to Realtime for infringement of the '530 patent pursuant to 35 U.S.C. § 271.

329. As a result of Rackspace's infringement of the '530 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Rackspace's infringement, but in no event less than a reasonable royalty for the use made of the invention by Rackspace, together with interest and costs as fixed by the Court.

SolidFire

330. On information and belief, NetApp and SolidFire (hereinafter,

"NetApp/SolidFire") have made, used, offered for sale, sold and/or imported into the United States NetApp/SolidFire products that infringe the '530 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp/SolidFire's compression products and services, such as, *e.g.*, the SolidFire all-flash storage system and all versions and variations thereof since the issuance of the '530 patent ("Accused Instrumentality").

331. On information and belief, NetApp/SolidFire has directly infringed and continues to infringe at least Claim 1 of the '530 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block. Upon information and belief, NetApp/SolidFire uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp/SolidFire's customers.

332. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "A system comprising: a memory device." *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire has architected its entire storage system accordingly, minimizing writes by

compressing and deduplicating data before writing to its flash SSDs. When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache.").

333. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream." *See, e.g.*, http://www.solidfire.com/platform/element-os/global-efficiencies:

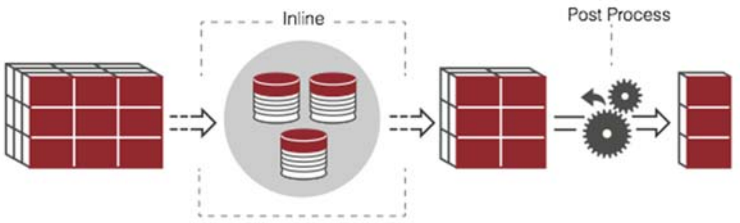

Two-layer compression

SolidFire utilizes both an in-line and post-process compression, optimizing for space without sacrificing performance.

334. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement "by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different". Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple

available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. An acknowledgement is returned to the host when — and only when — the data has been safely stored in the NVRAM of multiple storage nodes. Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash. If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated.")

335. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compressed data stream is stored on said memory device." *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks.

These blocks are immediately compressed and stored in the node's NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. … If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated.")

336. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form." *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system, and enables customers to: … Increase system performance by minimizing system resources … Improved performance is a key benefit of deduplication and compression techniques being performed inline as there is no performance tax that results from their usage within a primary storage infrastructure.").

337. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block." See, e.g., http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

338. On information and belief, NetApp/SolidFire also directly infringes and continues to infringe other claims of the '530 patent, for similar reasons as explained above with respect to

Claim 1 of the ‘530 patent.

339. On information and belief, all of the Accused Instrumentalities constitute the claimed systems in substantially the same way.

340. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the ‘530 patent.

341. On information and belief, NetApp/SolidFire has had knowledge of the ‘530 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘530 patent and knew of its infringement, including by way of this lawsuit.

342. NetApp/SolidFire’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the ‘530 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator, wherein said data accelerator is coupled to said memory device, a data stream is received by said data accelerator in received form, said data stream includes a first data block and a second data block, said data stream is compressed by said data accelerator to provide a compressed data stream by compressing said first data block with a first compression technique and said second data block with a second compression technique, said first and second compression techniques are different, said compressed data stream is stored on said memory device, said compression and storage occurs faster than said data stream is able to be stored on said memory device in said received form, a first data descriptor is stored on said memory device indicative of said first compression technique, and said first descriptor is utilized to decompress the portion of said compressed data stream associated with said first data block, thereby infringing the ‘530 patent. For example, NetApp/SolidFire instructs users of SolidFire about the benefits of its deduplication / compression features:

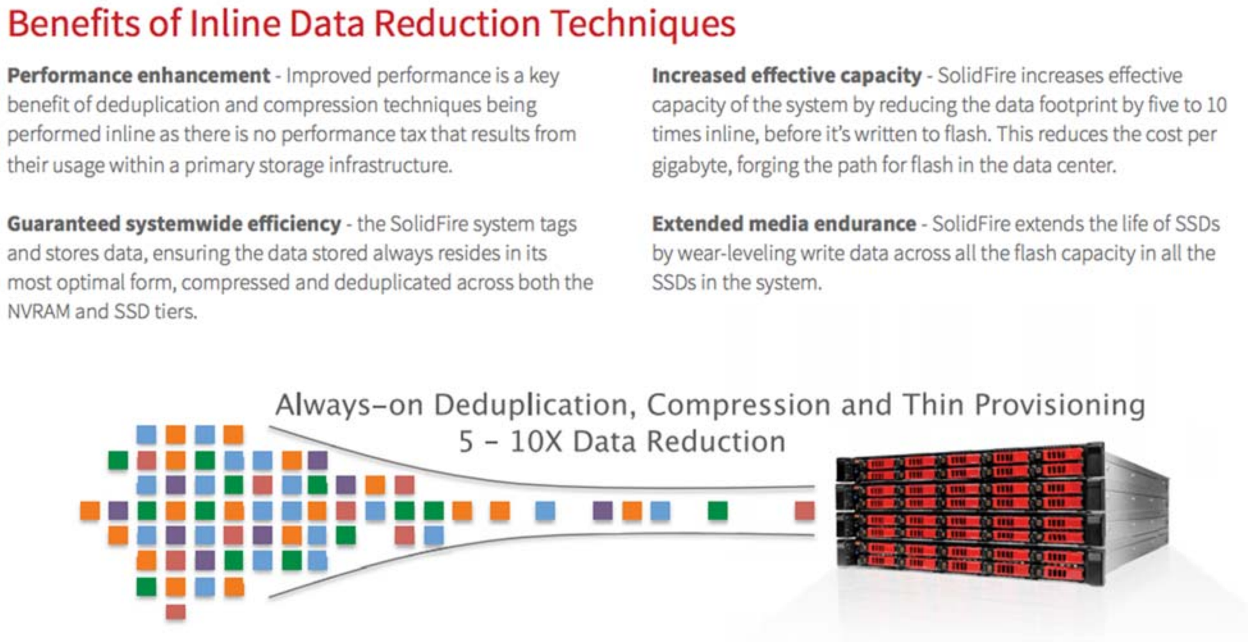

Benefits of Inline Data Reduction Techniques

Performance enhancement - Improved performance is a key benefit of deduplication and compression techniques being performed inline as there is no performance tax that results from their usage within a primary storage infrastructure.

Guaranteed systemwide efficiency - the SolidFire system tags and stores data, ensuring the data stored always resides in its most optimal form, compressed and deduplicated across both the NVRAM and SSD tiers.

Increased effective capacity - SolidFire increases effective capacity of the system by reducing the data footprint by five to 10 times inline, before it's written to flash. This reduces the cost per gigabyte, forging the path for flash in the data center.

Extended media endurance - SolidFire extends the life of SSDs by wear-leveling write data across all the flash capacity in all the SSDs in the system.

See also, e.g., http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. An acknowledgement is returned to the host when — and only when — the data has been safely stored in the NVRAM of multiple storage nodes. Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash. If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated."). Thus, with knowledge of the '530 patent gained from at least the filing and service of the original Complaint in this action, NetApp/SolidFire encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '530 patent, knowing that such use constituted infringement of the '530 patent.

343. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp/SolidFire has injured Realtime and is liable to Realtime for infringement of the '530 patent pursuant to 35 U.S.C. § 271.

344. As a result of NetApp/SolidFire's infringement of the '530 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp/SolidFire's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp/SolidFire, together with interest and costs as fixed by the Court.

COUNT VI

INFRINGEMENT OF U.S. PATENT NO. 8,643,513

345. Plaintiff realleges and incorporates by reference paragraphs 1-344 above, as if fully set forth herein.

346. Plaintiff Realtime is the owner by assignment of United States Patent No. 8,643,513 ("the '513 patent") entitled "Data compression systems and methods." The '513 patent was duly and legally issued by the United States Patent and Trademark Office on February 4, 2014. A true and correct copy of the '513 patent is included as Exhibit F.

NetApp ONTAP

347. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products and services that infringe the ‘513 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, NetApp’s compression products and services, such as, *e.g.*, NetApp ONTAP v8 (since November 2015) and ONTAP 9 (collectively, “ONTAP”) and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems (including but not limited to the NetApp AFF8000 series, including but not limited to the AFF8080 EX, AFF8060, and AFF8040) and hybrid disk FAS systems (including but not limited to NetApp FAS2500 series hybrid storage arrays, including but not limited to the FAS2554, FAS2552, and FAS2520, and NetApp FAS8000 Series hybrid storage systems, including but not limited to FAS8020, FAS8040, FAS8060, and FAS8080 EX systems), FlexPod® (including but not limited to FlexPod Datacenter, FlexPod Express, and FlexPod Select), ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the ‘513 patent (collectively, “Accused Instrumentality”).

348. On information and belief, NetApp has directly infringed and continues to infringe Claim 1 of the ‘513 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the ‘513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a

descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. Upon information and belief, NetApp uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

349. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion." *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

350. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data

block when the characteristic, attribute, or parameter is identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

351. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter

excludes analyzing based only on the descriptor." *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

352. On information and belief, NetApp also directly infringes and continues to infringe other claims of the '513 patent, for similar reasons as explained above with respect to Claim 1 of the '513 patent.

353. On information and belief, all of the Accused Instrumentalities constitute the claimed systems in substantially the same way.

354. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '513 patent.

355. On information and belief, NetApp has had knowledge of the '513 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the '513 patent and knew of its infringement, including by way of this lawsuit.

356. NetApp's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce end-users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the '513 patent by practicing compression methods claimed by the '513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to

be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. For example, NetApp touts the space-saving benefits of its inline data reduction technologies, including compression, deduplication, and compaction, to users of ONTAP. *See*, *e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers."). Thus, with knowledge of the '513 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '513 patent, knowing that such use constituted infringement of the '513 patent.

357. By making, using, offering for sale, selling and/or importing into the United States

the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the '513 patent pursuant to 35 U.S.C. § 271.

358. As a result of NetApp's infringement of the '513 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

NetApp AltaVault / RackSpace Private Cloud Powered By NetApp

359. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the '513 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the NetApp AltaVault product, NetApp ONTAP v8 (since November 2015) and ONTAP 9 (collectively, "ONTAP") and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems and hybrid disk FAS systems, FlexPod®, ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the '513 patent. Upon information and belief, RackSpace has collaborated with NetApp to make AltaVault compatible with RackSpace's cloud services. *See*, *e.g.*, http://solutionconnection.netapp.com/altavault ("AltaVault seamlessly integrates with an organization's current backup software, and with all leading public and private cloud storage providers. AltaVault partners fall into three categories: … Public cloud providers … Rackspace.").

360. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace and NetApp combination products and/or services that infringe the '513 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, "Rackspace Private Cloud Powered by NetApp", which integrates NetApp's deduplication and compression features into Rackspace Private Cloud. *See*, *e.g.*, http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp;

http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**"). Upon information and belief, the deduplication / compression features enabled by NetApp in "Rackspace Private Cloud Powered by NetApp" function in a similar manner as the deduplication / compression features in NetApp AltaVault, and both AltaVault and "Rackspace Private Cloud Powered by NetApp", and all versions and variations thereof since the issuance of the '513 patent (collectively, "Accused Instrumentality") infringe the '513 patent.

361. On information and belief, NetApp has directly infringed and continues to infringe Claim 1 of the '513 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the '513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based

only on the descriptor. Upon information and belief, NetApp uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

362. On information and belief, Rackspace has directly infringed and continues to infringe Claim 1 of the '513 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the '513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. Upon information and belief, Rackspace uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

363. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the

appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion." *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no compression: Disables compression of any data written to the share. This is useful if you are copying over already-compressed data (for example: photos, videos, or proprietary formats such as medical data that might be compressed and encrypted already)."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

364. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the

data block to provide a compressed data block when the characteristic, attribute, or parameter is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no-dedup: Specifies that data written to this share should not be checked for duplication. The AltaVault does not check if there is duplication of the data written to the share and not does perform de-duplication."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.")

365. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes

analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor." *See*, *e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Industry-leading data reduction. AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. Network and cloud optimization. Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."); https://library.netapp.com/ecm/ecm_get_file/ECMP12434740 at 151 ("no compression: Disables compression of any data written to the share. This is useful if you are copying over already-compressed data (for example: photos, videos, or proprietary formats such as medical data that might be compressed and encrypted already)."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

366. On information and belief, NetApp and Rackspace also directly infringe and continue to infringe other claims of the '513 patent, for similar reasons as explained above with respect to Claim 1 of the '513 patent.

367. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

368. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '513 patent.

369. On information and belief, NetApp and Rackspace have had knowledge of the '513 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, NetApp and Rackspace knew of the '513 patent and knew of their infringement, including by way of this lawsuit.

370. NetApp's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce end-users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the '513 patent by practicing compression methods claimed by the '513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. For example, NetApp emphasizes the benefits of the inline deduplication and compression features of NetApp AltaVault: "AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. … Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x."). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf. NetApp specifically intended and was aware that these

normal and customary activities would infringe the ‘513 patent. Thus, with knowledge of the ‘513 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘513 patent, knowing that such use constituted infringement of the ‘513 patent.

371. Rackspace’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce end-users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the ‘513 patent by practicing compression methods claimed by the ‘513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. For example, Rackspace instructs customers of the benefits of integrating NetApp storage solutions into Rackspace Private Cloud footprints: “Rackspace Private Cloud (RPC) powered by OpenStack … [can] connect your OpenStack cloud directly to a shared storage solution via Cinder integration drivers. This is where our friends at NetApp come into play. Rackspace and NetApp have formed a unique relationship to improve the Cinder shared storage capability within OpenStack. These two teams worked

together to create a repeatable, approved, and tested process to integrate NetApp storage solutions into Rackspace Private Cloud footprints within a Rackspace datacenter or at the customer's datacenter." *See* https://developer.rackspace.com/blog/rpc-and-netapp-hearts-block-storage/. NetApp storage solutions used with Rackspace Private Cloud powered by NetApp include infringing deduplication / compression features. *See, e.g.*, http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features."). Rackspace specifically intended and was aware that these normal and customary activities would infringe the '513 patent. Thus, with knowledge of the '513 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '513 patent, knowing that such use constituted infringement of the '513 patent.

372. For similar reasons, NetApp and Rackspace also induce their customers to use the Accused Instrumentalities to infringe other claims of the '513 patent. NetApp and Rackspace specifically intended and were aware that these normal and customary activities would infringe the '513 patent. NetApp and Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '513 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp and Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp and Rackspace have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '513 patent, knowing that

such use constitutes infringement of the ‘513 patent.

373. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp and Rackspace have injured Realtime and are liable to Realtime for infringement of the ‘513 patent pursuant to 35 U.S.C. § 271.

374. As a result of NetApp's and Rackspace's infringement of the ‘513 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's and Rackspace's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp and Rackspace, together with interest and costs as fixed by the Court.

NetApp Mars OS

375. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the ‘513 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the Mars OS software running on the FlashRay system and all versions and variations thereof since the issuance of the ‘513 patent ("Accused Instrumentality").

376. On information and belief, NetApp has directly infringed and continues to infringe Claim 1 of the ‘513 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the ‘513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is

identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. Upon information and belief, NetApp uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

377. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion." *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. … Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.")

378. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified". Even if the deduplication

function in the Accused Instrumentality were found to not literally meet the "analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.").

379. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor." *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: … Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.").

380. On information and belief, NetApp also directly infringes and continues to infringe other claims of the ‘513 patent, for similar reasons as explained above with respect to Claim 1 of the ‘513 patent.

381. On information and belief, all of the Accused Instrumentalities constitute the claimed systems in substantially the same way.

382. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the ‘513 patent.

383. On information and belief, NetApp has had knowledge of the ‘513 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘513 patent and knew of its infringement, including by way of this lawsuit.

384. NetApp’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce end-users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the ‘513 patent by practicing compression methods claimed by the ‘513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. For example, NetApp

instructs users of Mars OS about the benefits of its deduplication / compression features: "Industry-leading storage efficiency: Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs." *See* https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm. Thus, with knowledge of the '513 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '513 patent, knowing that such use constituted infringement of the '513 patent.

385. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the '513 patent pursuant to 35 U.S.C. § 271.

386. As a result of NetApp's infringement of the '513 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

Rackspace Cloud Backup

387. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace products that infringe the '513 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Rackspace's compression products and services, such as, e.g., Rackspace Cloud Backup and all

versions and variations thereof since the issuance of the '513 patent ("Accused Instrumentality").

388. On information and belief, Rackspace has directly infringed and continues to infringe Claim 1 of the '513 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the '513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. Upon information and belief, Rackspace uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

389. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion." *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-

level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip.”); https://www.rackspace.com/en-us/cloud/backup (“Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files.”); http://es.rack.ly/cloud/backup/how-it-works (“To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.”).

390. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified”. Deduplication literally meets a correct construction of “compression” – “[Representing / represented / representation] of data with fewer bits”. See Dkt. No. 362 in Case No. 6:15-cv-00463-RWS-JDL. Even if the deduplication function in the Accused Instrumentality were found to not literally meet the “analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent

algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data.").

391. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter

excludes analyzing based only on the descriptor." *See, e.g.,* https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

392. On information and belief, Rackspace also directly infringes and continues to infringe other claims of the '513 patent, for similar reasons as explained above with respect to Claim 1 of the '513 patent.

393. On information and belief, all of the Accused Instrumentalities operate in substantially the same way.

394. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '513 patent.

395. On information and belief, Rackspace has had knowledge of the '513 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief,

Rackspace knew of the ‘513 patent and knew of its infringement, including by way of this lawsuit.

396. Rackspace’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce end-users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the ‘513 patent by practicing compression methods claimed by the ‘513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. For example, Rackspace instructs users of Rackspace Cloud Backup about the benefits of its deduplication / compression features: “Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. Another benefit is that by using this method, you can retrieve previous versions of files, up to the limits specified by the customer-defined retention settings. De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a “delta” of the previous backup. These “deltas” are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required. To save on additional

overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip." See https://community.rackspace.com/products/f/25/t/7263. Rackspace also emphasizes the reduction in storage costs and backup time that the use of its deduplication/compression technology provides. See https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."). Rackspace specifically intended and was aware that these normal and customary activities would infringe the '513 patent. Thus, with knowledge of the '513 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '513 patent, knowing that such use constituted infringement of the '513 patent.

397. For similar reasons, Rackspace also induces its customers to use the Accused Instrumentalities to infringe other claims of the '513 patent. Rackspace specifically intended and was aware that these normal and customary activities would infringe the '513 patent. Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '513 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Rackspace has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '513 patent, knowing that such use constitutes infringement of the '513 patent.

398. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities'

compression features, Rackspace has injured Realtime and is liable to Realtime for infringement of the ‘513 patent pursuant to 35 U.S.C. § 271.

399. As a result of Rackspace's infringement of the ‘513 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Rackspace's infringement, but in no event less than a reasonable royalty for the use made of the invention by Rackspace, together with interest and costs as fixed by the Court.

SolidFire

400. On information and belief, NetApp and SolidFire (hereinafter, "NetApp/SolidFire") have made, used, offered for sale, sold and/or imported into the United States NetApp/SolidFire products that infringe the ‘513 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp/SolidFire's compression products and services, such as, *e.g.*, the SolidFire all-flash storage system and all versions and variations thereof since the issuance of the ‘513 patent ("Accused Instrumentality").

401. On information and belief, NetApp/SolidFire has directly infringed and continues to infringe Claim 1 of the ‘513 patent, for example, through its own use and testing of the Accused Instrumentalities to practice compression methods claimed by the ‘513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing

the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. Upon information and belief, NetApp/SolidFire uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp/SolidFire's customers.

402. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion." *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire has architected its entire storage system accordingly, minimizing writes by compressing and deduplicating data before writing to its flash SSDs. When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. … If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value.").

403. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified". Even if the deduplication function in the Accused Instrumentality were found to not literally meet the "analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the

data block to provide a compressed data block when the characteristic, attribute, or parameter is identified" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, deduplication performs substantially the same function (for example, reducing the overall amount of bits to store) in substantially the same way (by, for example, applying a technique based on the specific content of the incoming data in order to present for storage fewer overall bits) to achieve substantially the same result (for example, storage of fewer bits of data overall). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash.").

404. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor." *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically

identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … SolidFire has architected its entire storage system accordingly, minimizing writes by compressing and deduplicating data before writing to its flash SSDs. When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. … Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash. If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated.").

405. On information and belief, NetApp/SolidFire also directly infringes and continues to infringe other claims of the '513 patent, for similar reasons as explained above with respect to Claim 1 of the '513 patent.

406. On information and belief, all of the Accused Instrumentalities constitute the claimed systems in substantially the same way.

407. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '513 patent.

408. On information and belief, NetApp/SolidFire has had knowledge of the '513 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, NetApp/SolidFire knew of the '513 patent and knew of its infringement, including by way of this lawsuit.

409. NetApp/SolidFire's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce end-users of

the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the ‘513 patent by practicing compression methods claimed by the ‘513 patent, including a method of compressing a plurality of data blocks, comprising: analyzing the plurality of data blocks to recognize when an appropriate content independent compression algorithm is to be applied to the plurality of data blocks; applying the appropriate content independent data compression algorithm to a portion of the plurality of data blocks to provide a compressed data portion; analyzing a data block from another portion of the plurality of data blocks for recognition of any characteristic, attribute, or parameter that is indicative of an appropriate content dependent algorithm to apply to the data block; and applying the appropriate content dependent data compression algorithm to the data block to provide a compressed data block when the characteristic, attribute, or parameter is identified, wherein the analyzing the plurality of data blocks to recognize when the appropriate content independent compression algorithm is to be applied excludes analyzing based only on a descriptor indicative of the any characteristic, attribute, or parameter, and wherein the analyzing the data block to recognize the any characteristic, attribute, or parameter excludes analyzing based only on the descriptor. For example, NetApp/SolidFire instructs users of SolidFire about the benefits of its deduplication / compression features:

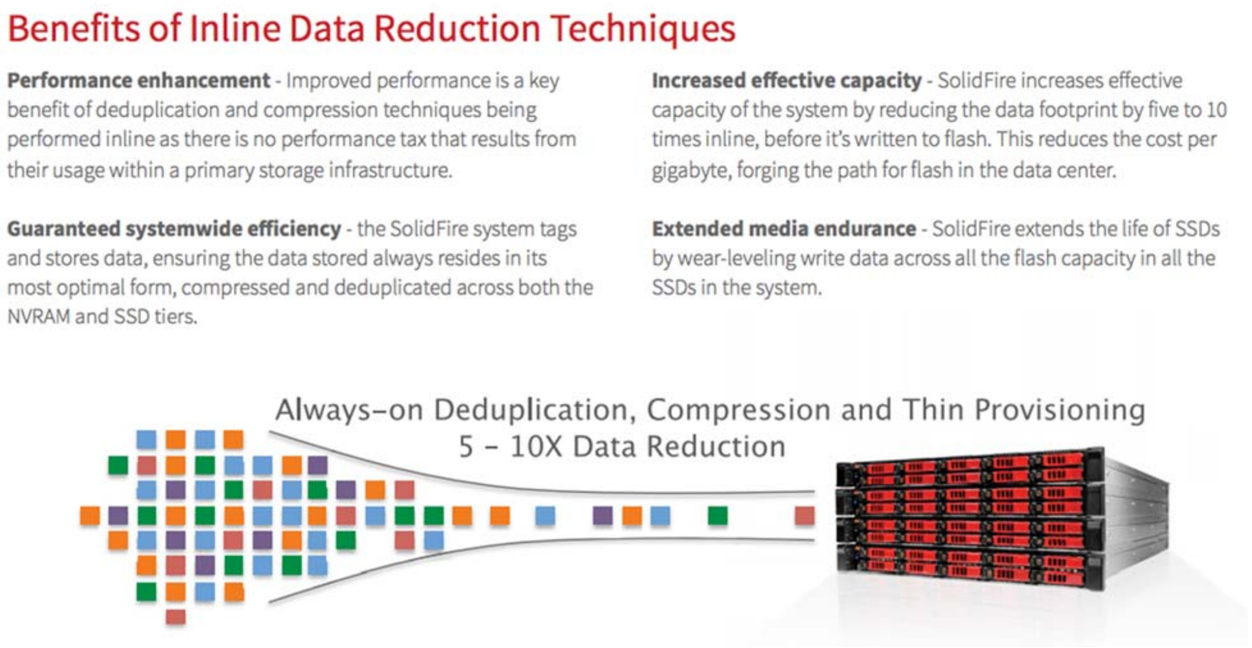

Benefits of Inline Data Reduction Techniques

Performance enhancement - Improved performance is a key benefit of deduplication and compression techniques being performed inline as there is no performance tax that results from their usage within a primary storage infrastructure.

Guaranteed systemwide efficiency - the SolidFire system tags and stores data, ensuring the data stored always resides in its most optimal form, compressed and deduplicated across both the NVRAM and SSD tiers.

Increased effective capacity - SolidFire increases effective capacity of the system by reducing the data footprint by five to 10 times inline, before it's written to flash. This reduces the cost per gigabyte, forging the path for flash in the data center.

Extended media endurance - SolidFire extends the life of SSDs by wear-leveling write data across all the flash capacity in all the SSDs in the system.

Always–on Deduplication, Compression and Thin Provisioning

5 – 10X Data Reduction

See also, e.g., http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. An acknowledgement is returned to the host when — and only when — the data has been safely stored in the NVRAM of multiple storage nodes. Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash. If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated."). Thus, with knowledge of the '513 patent gained from at least the filing and service of the original Complaint in this action, NetApp/SolidFire encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '513 patent, knowing that such use constituted infringement of the '513 patent.

410. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp/SolidFire has injured Realtime and is liable to Realtime for infringement of the '513 patent pursuant to 35 U.S.C. § 271.

411. As a result of NetApp/SolidFire's infringement of the '513 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp/SolidFire, together with interest and costs as fixed by the Court.

COUNT VII
INFRINGEMENT OF U.S. PATENT NO. 9,116,908

412. Plaintiff Realtime realleges and incorporates by reference paragraphs 1-411 above, as if fully set forth herein.

413. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,116,908 ("the '908 Patent") entitled "System and methods for accelerated data storage and retrieval." The '908 Patent was duly and legally issued by the United States Patent and Trademark Office on August 25, 2015. A true and correct copy of the '908 Patent is included as Exhibit G.

NetApp ONTAP

414. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products and services that infringe the '908 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, NetApp's compression products and services, such as, *e.g.*, NetApp ONTAP

v8 (since November 2015) and ONTAP 9 (collectively, "ONTAP") and all products and services running ONTAP, including but not limited to NetApp all-flash AFF systems (including but not limited to the NetApp AFF8000 series, including but not limited to the AFF8080 EX, AFF8060, and AFF8040) and hybrid disk FAS systems (including but not limited to NetApp FAS2500 series hybrid storage arrays, including but not limited to the FAS2554, FAS2552, and FAS2520, and NetApp FAS8000 Series hybrid storage systems, including but not limited to FAS8020, FAS8040, FAS8060, and FAS8080 EX systems), FlexPod® (including but not limited to FlexPod Datacenter, FlexPod Express, and FlexPod Select), ONTAP Select, FlexArray®, NetApp Private Storage (NPS) for Cloud, and ONTAP Cloud, and all versions and variations thereof since the issuance of the '908 patent (collectively, "Accused Instrumentality").

415. On information and belief, NetApp has directly infringed and continues to infringe the '908 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

416. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a system comprising: a memory device". *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative

inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

417. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See*, *e.g.*,

https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

418. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the compressed first and second data blocks are stored on the memory device". *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case."); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx ("Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.").

419. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form." *See, e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 ("All Flash FAS systems are built with

innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case.”); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx (“Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.”).

420. On information and belief, NetApp also directly infringed and continued to infringe other claims of the ‘908 patent, for similar reasons as explained above with respect to Claim 1 of the ‘908 patent.

421. On information and belief, all of the Accused Instrumentalities constitute the claimed system in substantially the same way.

422. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the ‘908 patent.

423. On information and belief, NetApp has had knowledge of the ‘908 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘908 patent and knew of its infringement, including by way of this lawsuit.

424. NetApp’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the ‘908 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed

first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the ‘908 Patent. For example, NetApp touts the space-saving benefits of its inline data reduction technologies, including compression, deduplication, and compaction, to users of ONTAP. *See*, *e.g.*, https://community.netapp.com/t5/Technology/Maximize-Performance-amp-Control-Costs-with-NetApp-Solutions-for-MongoDB/ba-p/119867 (“All Flash FAS systems are built with innovative inline data reduction technologies including inline adaptive compression, inline deduplication, and inline data compaction introduced in the latest ONTAP 9 release. These technologies provide space savings of 5 to 10 times, on average, for a typical use case.”); http://www.netapp.com/us/products/platform-os/ontap/storage-efficiency.aspx (“Reduce your data storage and management costs using NetApp® ONTAP® data management software and its leading set of optimization and data-reduction technologies. … Use inline data-reduction technologies (deduplication, compression, and compaction) across multiple applications and storage tiers.”). Thus, with knowledge of the ‘908 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘908 patent, knowing that such use constituted infringement of the ‘908 patent.

425. For similar reasons, NetApp also induces its customers to use the Accused Instrumentalities to infringe other claims of the ‘908 patent. NetApp specifically intended and was aware that these normal and customary activities would infringe the ‘908 patent. NetApp performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘908 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘908

patent, knowing that such use constitutes infringement of the ‘908 patent.

426. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the ‘908 patent pursuant to 35 U.S.C. § 271.

427. As a result of NetApp's infringement of the ‘908 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

NetApp AltaVault / RackSpace Private Cloud Powered By NetApp

428. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the ‘908 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the NetApp AltaVault product and all versions and variations thereof since the issuance of the ‘908 patent. Upon information and belief, RackSpace has collaborated with NetApp to make AltaVault compatible with RackSpace's cloud services. *See*, *e.g.*, http://solutionconnection.netapp.com/altavault ("AltaVault seamlessly integrates with an organization's current backup software, and with all leading public and private cloud storage providers. AltaVault partners fall into three categories: … Public cloud providers … Rackspace.").

429. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace and NetApp combination products and/or services that infringe the ‘908 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, "Rackspace Private Cloud Powered by NetApp", which integrates NetApp's deduplication and compression features into Rackspace Private Cloud. See, e.g., http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentI

D=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features."). Upon information and belief, the deduplication / compression features enabled by NetApp in "Rackspace Private Cloud Powered by NetApp" function in a similar manner as the deduplication / compression features in NetApp AltaVault, and both AltaVault and "Rackspace Private Cloud Powered by NetApp", and all versions and variations thereof since the issuance of the '908 patent (collectively, "Accused Instrumentality") infringe the '908 patent.

430. On information and belief, NetApp has directly infringed and continues to infringe the '908 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

431. On information and belief, Rackspace has directly infringed and continues to infringe the '908 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused

Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. Upon information and belief, Rackspace uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

432. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a system comprising: a memory device". *See, e.g.*, http://www.netapp.com/us/media/ds-3687.pdf ("AltaVault physical appliances are the industry's most scalable cloud-integrated storage appliances, with capacities ranging from 32TB up to 384TB of usable local cache. … AltaVault physical appliances are built on a scalable and efficient hardware platform that is optimized to reduce data footprints and rapidly stream data to the cloud."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

433. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a data accelerator configured to compress: (i) a first data block

with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf:

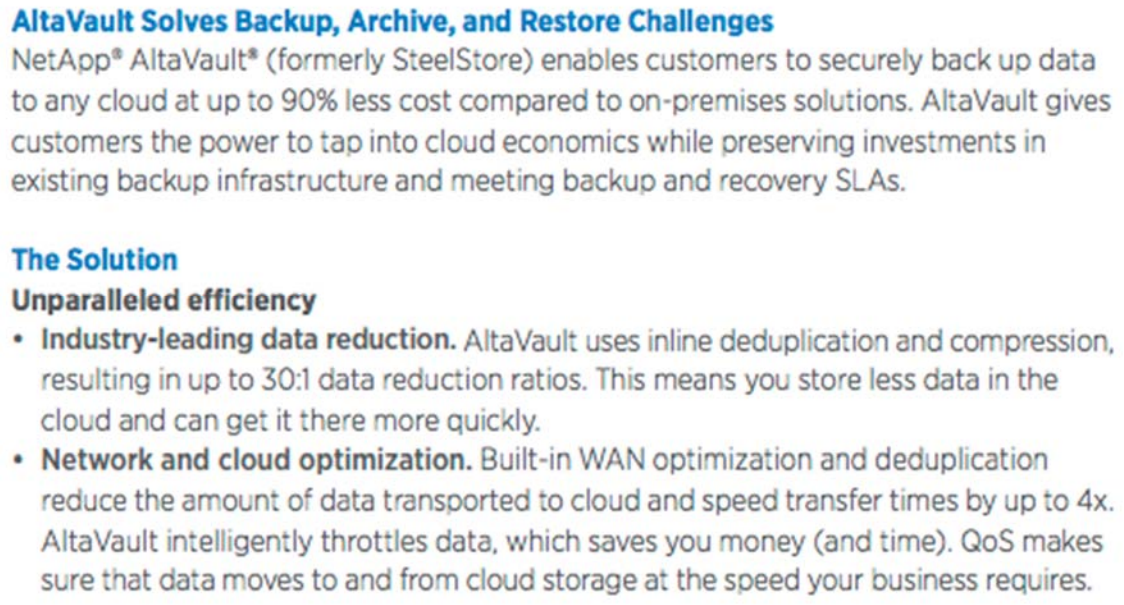

AltaVault Solves Backup, Archive, and Restore Challenges

NetApp® AltaVault® (formerly SteelStore) enables customers to securely back up data to any cloud at up to 90% less cost compared to on-premises solutions. AltaVault gives customers the power to tap into cloud economics while preserving investments in existing backup infrastructure and meeting backup and recovery SLAs.

The Solution

Unparalleled efficiency

- **Industry-leading data reduction.** AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly.
- **Network and cloud optimization.** Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x. AltaVault intelligently throttles data, which saves you money (and time). QoS makes sure that data moves to and from cloud storage at the speed your business requires.

http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's storage infrastructure enables the storage features.**").

434. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the compressed first and second data blocks are stored on the memory device". *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf:

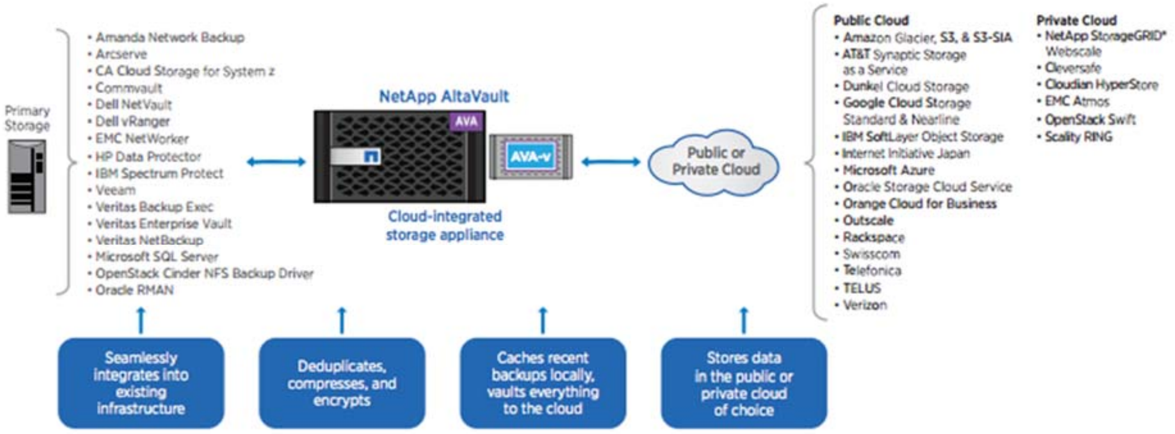
Primary Storage
• Amanda Network Backup
• Arcserve
• CA Cloud Storage for System z
• Commvault
• Dell NetVault
• Dell vRanger
• EMC NetWorker
• HP Data Protector
• IBM Spectrum Protect
• Veeam
• Veritas Backup Exec
• Veritas Enterprise Vault
• Veritas NetBackup
• Microsoft SQL Server
• OpenStack Cinder NFS Backup Driver
• Oracle RMAN
NetApp AltaVault
AVA
AVA-v
Cloud-integrated storage appliance
Public or Private Cloud
Public Cloud
• Amazon Glacier, S3, & S3-SIA
• AT&T Synaptic Storage as a Service
• Dunkel Cloud Storage
• Google Cloud Storage Standard & Nearline
• IBM SoftLayer Object Storage
• Internet Initiative Japan
• Microsoft Azure
• Oracle Storage Cloud Service
• Orange Cloud for Business
• Outscale
• Rackspace
• Swisscom
• Telefonica
• TELUS
• Verizon
Private Cloud
• NetApp StorageGRID* Webscale
• Cleversafe
• Cloudian HyperStore
• EMC Atmos
• OpenStack Swift
• Scality RING
Seamlessly integrates into existing infrastructure
Deduplicates, compresses, and encrypts
Caches recent backups locally, vaults everything to the cloud
Stores data in the public or private cloud of choice

Figure 1) NetApp AltaVault provides seamless integrations with existing applications and cloud service providers.

http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features.").

435. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form." *See*, *e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf ("Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x. AltaVault intelligently throttles data, which saves you money (and time)."); http://solutionconnection.netapp.com/rackspace-private-cloud-powered-by-netapp; http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, **NetApp enables features like deduplication, compression**, and NetApp SnapMirror and FlexClone technology **within the context of a Rackspace Private Cloud user experience**. Rackspace enables the storage drivers, **NetApp's**

storage infrastructure enables the storage features.”).

436. On information and belief, NetApp and Rackspace also directly infringe and continue to infringe other claims of the ‘908 patent, for similar reasons as explained above with respect to Claim 1 of the ‘908 patent.

437. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

438. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the ‘908 patent.

439. On information and belief, NetApp and RackSpace have had knowledge of the ‘908 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp and Rackspace knew of the ‘908 patent and knew of their infringement, including by way of this lawsuit.

440. NetApp’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the ‘908 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the ‘908 Patent. For example, NetApp emphasizes the benefits of the inline deduplication and compression features of NetApp AltaVault to its customers: “AltaVault uses inline deduplication and compression, resulting in up to 30:1 data reduction ratios. This means you store less data in the cloud and can get it there more quickly. …

Built-in WAN optimization and deduplication reduce the amount of data transported to cloud and speed transfer times by up to 4x.”). *See, e.g.*, http://www.scality.com/wp-content/uploads/2016/06/NetApp-AltaVault-Cloud-Integrated-Storage.pdf. NetApp also explains that its deduplication/compression approach is superior to others and results in faster restore times: “When you back up to AltaVault, it performs inline (real-time) deduplication of the backup data and replicates data into the cloud. AltaVault uses the local disk to store enough data for recovery of most recent backups. Such a mechanism provides LAN performance for the most likely restores. This deduplication process uses variable segment length inline deduplication plus compression, which is superior to other techniques such as fixed block. AltaVault deduplication level typically ranges between 10 and 30x. Deduplication performance depends on the incoming data type so turn off encryption and compression in the backup applications. Use the native encryption and deduplication in AltaVault to get higher data reduction rates than other typical software products. … AltaVault also optimizes restores from the cloud because it recalls only deduplicated data (which is not in the local cache) from the cloud. So if the customer is getting 10x deduplication, for example, and he or she needs to restore 10 TB of data, AltaVault needs only about 1 TB to restore. Over a 100-Mb line, this results in a time saving of days.” See https://library.netapp.com/ecm/ecm_download_file/ECMP12434738. Thus, with knowledge of the ‘908 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘908 patent, knowing that such use constituted infringement of the ‘908 patent.

441. Rackspace’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the ‘908 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator configured

to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. For example, Rackspace instructs customers of the benefits of integrating NetApp storage solutions into Rackspace Private Cloud footprints: "Rackspace Private Cloud (RPC) powered by OpenStack … [can] connect your OpenStack cloud directly to a shared storage solution via Cinder integration drivers. This is where our friends at NetApp come into play. Rackspace and NetApp have formed a unique relationship to improve the Cinder shared storage capability within OpenStack. These two teams worked together to create a repeatable, approved, and tested process to integrate NetApp storage solutions into Rackspace Private Cloud footprints within a Rackspace datacenter or at the customer's datacenter." *See* https://developer.rackspace.com/blog/rpc-and-netapp-hearts-block-storage/. NetApp storage solutions used with Rackspace Private Cloud powered by NetApp include infringing deduplication / compression features. *See, e.g.,* http://solutionconnection.netapp.com/Core/DownloadDoc.aspx?documentID=125690&contentID=236630 ("Rackspace enables you to use NetApp drivers to seamlessly connect your NetApp storage to the Rackspace Private Cloud interface that sits on top of your architecture. Delivering maximum efficiency and compatibility, NetApp enables features like deduplication, compression, and NetApp SnapMirror and FlexClone technology within the context of a Rackspace Private Cloud user experience. Rackspace enables the storage drivers, NetApp's storage infrastructure enables the storage features."). Thus, with knowledge of the '908 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '908 patent, knowing that such use constituted infringement of the '908 patent.

442. For similar reasons, NetApp and Rackspace also induce their customers to use the

Accused Instrumentalities to infringe other claims of the '908 patent. NetApp and Rackspace specifically intended and were aware that these normal and customary activities would infringe the '908 patent. NetApp and Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '908 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp and Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp and Rackspace have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '908 patent, knowing that such use constitutes infringement of the '908 patent.

443. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp and Rackspace have injured Realtime and are liable to Realtime for infringement of the '908 patent pursuant to 35 U.S.C. § 271.

444. As a result of NetApp's and Rackspace's infringement of the '908 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's and Rackspace's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp and Rackspace, together with interest and costs as fixed by the Court.

NetApp Mars OS

445. On information and belief, NetApp has made, used, offered for sale, sold and/or imported into the United States NetApp products that infringe the '908 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp's compression products and services, such as, *e.g.*, the Mars OS software running on the FlashRay system and all versions and variations thereof since the issuance of the '908 patent ("Accused Instrumentality").

446. On information and belief, NetApp has directly infringed and continues to infringe the '908 patent, for example, through its own use, testing, sale, offer for sale, and/or importation

of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. Upon information and belief, NetApp uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp's customers.

447. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a system comprising: a memory device". *See, e.g.*, http://www.netapp.com/us/company/news/press-releases/news-rel-20140917-816753.aspx ("NetApp (NASDAQ: NTAP) today announced its first shipments of FlashRay™, the company's purpose-built all-flash storage array with the new NetApp® Mars™ operating system. FlashRay is designed from the ground up to improve the performance, efficiency, and manageability of all-flash storage architectures used in enterprise application environments."); https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS is the software running on the FlashRay system, which stores data on an all-SSD (solid-state drive) shelf.").

448. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously

compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.").

449. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the compressed first and second data blocks are stored on the memory device". *See, e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm ("Mars OS is the software running on the FlashRay system, which stores data on an all-SSD (solid-

state drive) shelf. … Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings.”).

450. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form.” *See*, *e.g.*, https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm (“Mars OS delivers high performance in the form of consistent and predictable low latency and high throughput. For small random access workloads, throughput is measured in I/O operations per second (IOPS); for sequential and large random workloads, it is measured as bandwidth (for example, GB/s).”).

451. On information and belief, NetApp also directly infringed and continued to infringe other claims of the ‘908 patent, for similar reasons as explained above with respect to Claim 1 of the ‘908 patent.

452. On information and belief, all of the Accused Instrumentalities constitute the claimed system in substantially the same way.

453. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the ‘908 patent.

454. On information and belief, NetApp has had knowledge of the ‘908 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp knew of the ‘908 patent and knew of its infringement, including by way of this lawsuit.

455. NetApp’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the ‘908 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are

converted into infringing systems comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. For example, NetApp instructs users of Mars OS about the benefits of its deduplication / compression features: "Industry-leading storage efficiency: Mars OS offers these features to maximize efficiency: Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings. Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system. Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs." *See* https://library.netapp.com/ecmdocs/ECMP12394203/html/Overview-Mgmt/features-mars.htm. Thus, with knowledge of the '908 patent gained from at least the filing and service of the original Complaint in this action, NetApp encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '908 patent, knowing that such use constituted infringement of the '908 patent.

456. For similar reasons, NetApp also induces its customers to use the Accused Instrumentalities to infringe other claims of the '908 patent. NetApp specifically intended and was aware that these normal and customary activities would infringe the '908 patent. NetApp performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '908 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief,

NetApp engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '908 patent, knowing that such use constitutes infringement of the '908 patent.

457. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp has injured Realtime and is liable to Realtime for infringement of the '908 patent pursuant to 35 U.S.C. § 271.

458. As a result of NetApp's infringement of the '908 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp, together with interest and costs as fixed by the Court.

Rackspace Cloud Backup

459. On information and belief, Rackspace has made, used, offered for sale, sold and/or imported into the United States Rackspace products that infringe the '908 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Rackspace's compression products and services, such as, e.g., Rackspace Cloud Backup and all versions and variations thereof since the issuance of the '908 patent ("Accused Instrumentality").

460. On information and belief, Rackspace has directly infringed and continues to infringe the '908 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data

blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. Upon information and belief, Rackspace uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to Rackspace's customers.

461. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a system comprising: a memory device". *See, e.g.*, https://www.rackspace.com/cloud/servers/features ("Cloud Servers reside in our world-class data centers, with ECC memory, and fully redundant networking and power all the way to the host. All virtual servers are backed by hardware RAID 10 storage.").

462. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block." Deduplication literally meets a correct construction of "compression" – "[Representing / represented / representation] of data with fewer bits". See Dkt. No. 362 in Case No. 6:15-cv-00463-RWS-JDL. Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of

the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.,* https://community.rackspace.com/products/f/25/t/7263 ("Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. … To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip."); https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."); http://es.rack.ly/cloud/backup/how-it-works ("To improve backup times and reduce size and storage costs, Rackspace Cloud Backup includes file compression and block-level de-duplication. The agent will know previous backups and utilises de-duplication technology to only back up changes and new additions, saving time by not replicating previously backed up data. It will then compress the new backup data, and apply the optional AES-256 encryption before any of the data leaves your server. Once the incremental backup is completed, we may additionally compress the files if it reduces the overall size.").

463. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the compressed first and second data blocks are stored on the memory device". After deduplication and compression, the compressed data blocks are stored on the memory device of the Accused Instrumentality. *See* https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental

to lower your total costs.”).

464. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, “and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form.” *See, e.g.*, https://community.rackspace.com/products/f/25/t/7263 (“De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a “delta” of the previous backup. These “deltas” are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required.”); https://www.rackspace.com/en-gb/cloud/backup/features (“Cloud Backup includes file compression and block-level deduplication to reduce backup times, size, and storage costs. After the initial backup of a data set, block-level deduplication saves time and storage space by only backing up files that have changed since the last backup.”).

465. On information and belief, Rackspace also directly infringes and continues to infringe other claims of the ‘908 patent, for similar reasons as explained above with respect to Claim 1 of the ‘908 patent.

466. On information and belief, all of the Accused Instrumentalities constitute the claimed system in substantially the same way.

467. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the ‘908 patent.

468. On information and belief, RackSpace has had knowledge of the ‘908 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, Rackspace knew of the ‘908 patent and knew of its infringement, including by way of this lawsuit.

469. Rackspace’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the ‘908 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are

converted into infringing systems comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. For example, Rackspace instructs users of Rackspace Cloud Backup about the benefits of its deduplication / compression features: "Rackspace Cloud Backup uses block-level de-duplication, which means we only save those parts of a file that have changed. In this way, a unique piece of data is saved only once, which maximizes the effectiveness of the backup, while minimizing your storage overhead. Another benefit is that by using this method, you can retrieve previous versions of files, up to the limits specified by the customer-defined retention settings. De-duplication also gives you flexibility because, with the possible exception of your first complete backup, every subsequent backup is just a "delta" of the previous backup. These "deltas" are smaller in size and overall make for faster backups and faster restores and also reduce the amount of storage required. To save on additional overhead, we may then compress the files if it reduces the size of the block. You can expect compression rates equivalent to those of gzip." See https://community.rackspace.com/products/f/25/t/7263. Rackspace also emphasizes the reduction in storage costs and backup time that the use of its deduplication/compression technology provides. See https://www.rackspace.com/en-us/cloud/backup ("Block-level compression and deduplication reduce storage costs by up to 20x compared to uncompressed backups. After your first backup, additional backups are incremental to lower your total costs. … Cloud Backup includes file compression and block-level deduplication to reduce backup times, file size, and storage costs. After the initial backup, block-level deduplication saves time and storage space by only backing up changed files."). Thus, with knowledge of the '908 patent gained from at least the filing and service of the original Complaint in this action, Rackspace encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the ‘908 patent, knowing that such use constituted infringement of the ‘908 patent.

470. For similar reasons, Rackspace also induces its customers to use the Accused Instrumentalities to infringe other claims of the ‘908 patent. Rackspace specifically intended and was aware that these normal and customary activities would infringe the ‘908 patent. Rackspace performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ‘908 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Rackspace engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Rackspace has induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘908 patent, knowing that such use constitutes infringement of the ‘908 patent.

471. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ compression features, Rackspace has injured Realtime and is liable to Realtime for infringement of the ‘908 patent pursuant to 35 U.S.C. § 271.

472. As a result of Rackspace’s infringement of the ‘908 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Rackspace’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Rackspace, together with interest and costs as fixed by the Court.

SolidFire

473. On information and belief, NetApp and SolidFire (hereinafter, “NetApp/SolidFire”) have made, used, offered for sale, sold and/or imported into the United States NetApp/SolidFire products that infringe the ‘908 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, NetApp/SolidFire’s compression products and services, such as, *e.g.*, the SolidFire all-flash storage system and all versions and variations thereof since the issuance of the ‘908 patent (“Accused Instrumentality”).

474. On information and belief, NetApp/SolidFire has directly infringed and continues to infringe the '908 patent, for example, through its own use, testing, sale, offer for sale, and/or importation of the Accused Instrumentalities and computer systems running the Accused Instrumentalities, which when used as designed and intended, constitute a system comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the '908 Patent. Upon information and belief, NetApp/SolidFire uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support for the Accused Instrumentality to NetApp/SolidFire's customers.

475. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a system comprising: a memory device". *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire has architected its entire storage system accordingly, minimizing writes by compressing and deduplicating data before writing to its flash SSDs. When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache.").

476. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block." Even if the determination of whether particular data within a data block of an input data stream is duplicative of data that has been previously

compressed and/or stored by the Accused Instrumentality were found not to literally meet the "compression technique" limitation, this limitation is met under the doctrine of equivalents because it is insubstantially different from what the limitation literally requires. Moreover, determining whether particular data within a data block of an input data stream is duplicative of data that has been previously compressed and/or stored by the Accused Instrumentality performs substantially the same function (for example, to provide the Accused Instrumentality with some parameter of the data that can be used as a basis to select the optimal data compression method among multiple available data compression methods) in substantially the same way (by, for example, identifying some characteristic of the data, beyond a mere descriptor that is indicative of the data type of the data within the data block, that is relevant to selecting among multiple available data compression methods) to achieve substantially the same result (for example, enabling the Accused Instrumentality to select the optimal data compression method from among multiple available data compression methods). *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. An acknowledgement is returned to the host when — and only when — the data has been safely stored in the NVRAM of multiple storage nodes. Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash. If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated.").

477. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "wherein the compressed first and second data blocks are stored on the memory device". *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. … If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated.").

478. The Accused Instrumentality satisfies literally and/or under the doctrine of equivalents the claim requirement, "and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form." *See, e.g.*, http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system, and enables customers to: … Increase system performance by minimizing system resources … Improved performance is a key benefit of deduplication and compression techniques being performed inline as there is no performance tax that results from their usage within a primary storage infrastructure.").

479. On information and belief, NetApp/SolidFire also directly infringed and continued to infringe other claims of the '908 patent, for similar reasons as explained above with respect to Claim 1 of the '908 patent.

480. On information and belief, all of the Accused Instrumentalities constitute the

claimed system in substantially the same way.

481. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the ‘908 patent.

482. On information and belief, NetApp/SolidFire has had knowledge of the ‘908 patent at least since the filing of this Complaint or shortly thereafter, and on information and belief, NetApp/SolidFire knew of the ‘908 patent and knew of its infringement, including by way of this lawsuit.

483. NetApp/SolidFire’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way on compatible systems to infringe the ‘908 patent, knowing that when the Accused Instrumentalities are used in their ordinary and customary manner with such compatible systems, such systems are converted into infringing systems comprising: a memory device; and a data accelerator configured to compress: (i) a first data block with a first compression technique to provide a first compressed data block; and (ii) a second data block with a second compression technique, different from the first compression technique, to provide a second compressed data block; wherein the compressed first and second data blocks are stored on the memory device, and the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form, thereby infringing the ‘908 Patent. For example, NetApp/SolidFire instructs users of SolidFire about the benefits of its deduplication / compression features:

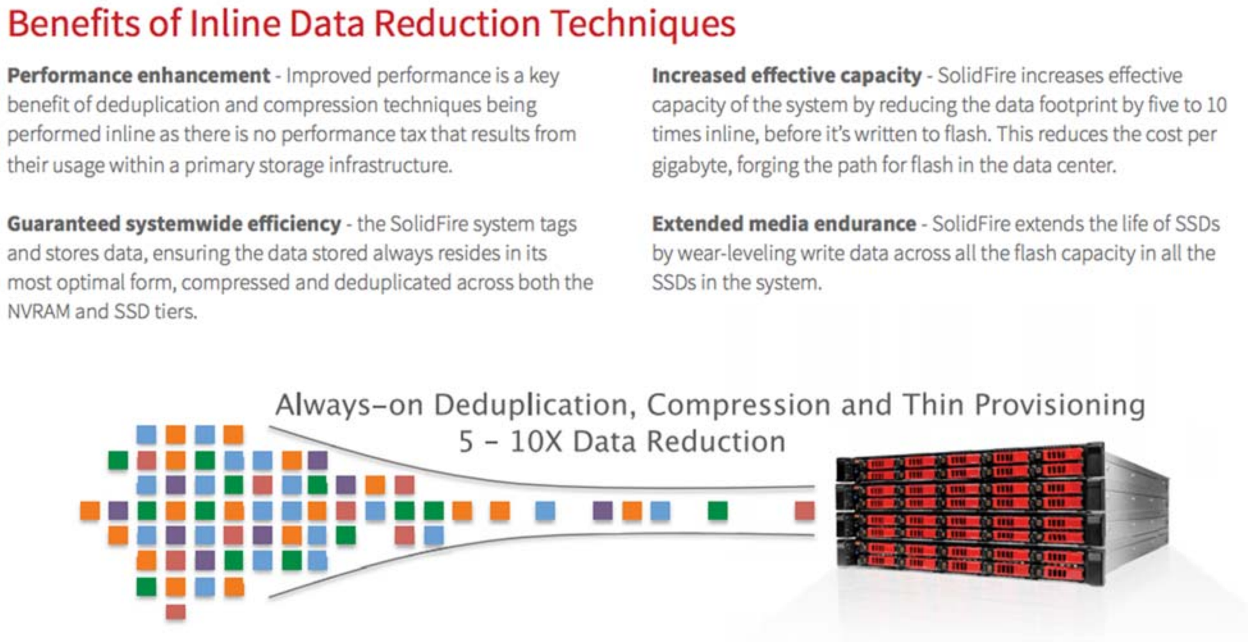

Benefits of Inline Data Reduction Techniques

Performance enhancement - Improved performance is a key benefit of deduplication and compression techniques being performed inline as there is no performance tax that results from their usage within a primary storage infrastructure.

Guaranteed systemwide efficiency - the SolidFire system tags and stores data, ensuring the data stored always resides in its most optimal form, compressed and deduplicated across both the NVRAM and SSD tiers.

Increased effective capacity - SolidFire increases effective capacity of the system by reducing the data footprint by five to 10 times inline, before it's written to flash. This reduces the cost per gigabyte, forging the path for flash in the data center.

Extended media endurance - SolidFire extends the life of SSDs by wear-leveling write data across all the flash capacity in all the SSDs in the system.

See also, e.g., http://info.solidfire.com/rs/538-SKP-058/images/SolidFire-Data-Efficiencies-Breif.pdf ("SolidFire's Deduplication Block Service receives a WriteBlock message, which contains the unique BlockID and data. The BlockID is then computed based on a hash from the 4K data. As a result, the Block Service will automatically identify whether the BlockID has already been written. If the Block Service recognizes the BlockID, duplicate objects will never translate into physical data writes onto the block drives. This entire process is performed in-line with no performance impact to the system … When a host writes data to a SolidFire storage node, that write is divided into 4KB data blocks. These blocks are immediately compressed and stored in the node's NVRAM write cache. Each compressed block is synchronously replicated to one or more additional storage nodes for data protection. An acknowledgement is returned to the host when — and only when — the data has been safely stored in the NVRAM of multiple storage nodes. Data contained in the compressed data chunk is then hashed. The system looks for that hash value in its index of stored data, which is distributed across the entire cluster. If the data is already present, the SolidFire operating system updates its metadata to indicate that the previously stored block should be delivered when the host reads the data being written, and the newly written block is discarded without ever being written to flash. If the compressed data block has unique data, it is stored in the system's block pool. The block pool is organized only by the content hash value, rather than by when data was written or from where it originated."). Thus, with knowledge of the '908 patent gained from at least the filing and service of the original Complaint in this action, NetApp/SolidFire encouraged users of the Accused Instrumentalities to use their deduplication/compression functionality to infringe the '908 patent, knowing that such use constituted infringement of the '908 patent.

484. For similar reasons, NetApp/SolidFire also induces its customers to use the Accused Instrumentalities to infringe other claims of the '908 patent. NetApp/SolidFire specifically intended and was aware that these normal and customary activities would infringe the '908 patent. NetApp/SolidFire performed the acts that constitute induced infringement, and would

induce actual infringement, with the knowledge of the ‘908 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, NetApp/SolidFire engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, NetApp/SolidFire has induced and continues to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the ‘908 patent, knowing that such use constitutes infringement of the ‘908 patent.

485. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, NetApp/SolidFire has injured Realtime and is liable to Realtime for infringement of the ‘908 patent pursuant to 35 U.S.C. § 271.

486. As a result of NetApp's infringement of the ‘908 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for NetApp/SolidFire's infringement, but in no event less than a reasonable royalty for the use made of the invention by NetApp/SolidFire, together with interest and costs as fixed by the Court.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff Realtime respectfully requests that this Court enter:

a. A judgment in favor of Plaintiff that NetApp, Rackspace, and SolidFire have infringed, either literally and/or under the doctrine of equivalents, the ‘506 patent, the ‘728 patent, the ‘992 patent, the ‘530 patent, the ‘513 patent, and the ‘908 patent; and that Rackspace has infringed the ‘867 patent;

b. A judgment and order requiring NetApp, Rackspace, and SolidFire to pay Plaintiff its damages, costs, expenses, and prejudgment and post-judgment interest for their infringement of the ‘506 patent, the ‘728 patent, the ‘992 patent, the ‘530 patent, the ‘513 patent, and the ‘908 patent; and a judgment and order requiring Rackspace to pay Plaintiff its damages, costs, expenses, and prejudgment and post-judgment interest for its infringement of the ‘867 patent as provided under 35 U.S.C. § 284;

c. A judgment and order requiring NetApp, Rackspace, and SolidFire to provide an accounting and to pay supplemental damages to Realtime, including without limitation, prejudgment and post-judgment interest;

d. A judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to Plaintiff its reasonable attorneys' fees against NetApp, Rackspace, and SolidFire; and

e. Any and all other relief as the Court may deem appropriate and just under the circumstances.

DEMAND FOR JURY TRIAL

Plaintiff, under Rule 38 of the Federal Rules of Civil Procedure, requests a trial by jury of any issues so triable by right.

Dated: October 11, 2016

Respectfully submitted,

/s/ Marc A. Fenster by permission Claire Henry
Marc A. Fenster (CA SBN 181067)
LEAD ATTORNEY
Reza Mirzaie (CA SBN 246953)
Brian D. Ledahl (CA SBN 186579)
Jeffrey Z.Y. Liao (CA SBN 288994)
C. Jay Chung (CA SBN 252794)
RUSS AUGUST & KABAT
12424 Wilshire Boulevard, 12th Floor
Los Angeles, CA 90025
(310) 826-7474
mfenster@raklaw.com
rmirzaie@raklaw.com
bledahl@raklaw.com
jliao@raklaw.com
jchung@raklaw.com

T. John Ward, Jr.
Texas State Bar No. 00794818
E-mail: jw@wsfirm.com
Claire Abernathy Henry
Texas State Bar No. 24053063
E-mail: claire@wsfirm.com
Andrea L. Fair

Texas State Bar No. 24078488
E-mail: andrea@wsfirm.com
WARD, SMITH & HILL, PLLC
PO Box 1231
Longview, Texas 75606
(903) 757-6400 (telephone)
(903) 757-2323 (facsimile)

Attorneys for Plaintiff
Realtime Data LLC d/b/a IXO

CERTIFICATE OF SERVICE

I hereby certify that a copy of the foregoing document was served on all counsel of record by email on this 11th day of October, 2016.

/s/ Claire Abernathy Henry