

Exhibit C5
Invalidity Chart:
U.S. Patent No. 7,152,182 (“’182 Patent”) (Ji et al.)
in view of
Managing MC/ServiceGuard (“ServiceGuard Manual”)

Managing MC/ServiceGuard, published in March 2002 by Hewlett Packard (“ServiceGuard Manual”) anticipates and/or renders obvious claims 1-7 and 9-18 (“Asserted Claims”) of the ’182 Patent under 35 U.S.C. §102 and/or §103, either alone or in combination with other prior art references, and/or in combination with the knowledge of a person of ordinary skill. ServiceGuard Manual anticipates and/or renders obvious the Asserted Claims of the ’182 Patent as those claims have been applied by Valtrus in its Infringement Contentions and/or under SAP’s understanding of the proper construction of the claims.

Nothing stated herein shall be treated as an admission or suggestion that SAP agrees with Valtrus regarding either the scope of any of the Asserted Claims or the claim constructions advanced by Valtrus in its Infringement Contentions or anywhere else, or that any of SAP’s accused products meet any limitation of the Asserted Claims. Nothing stated herein shall be construed as an admission or a waiver of any particular construction of any claim term. Nothing stated herein shall be construed as an admission that the Asserted Claims are directed to patent-eligible subject matter. SAP reserve all rights to challenge any of the Asserted Claims under 35 U.S.C. § 101. SAP also reserve all rights to challenge any of the claim terms herein under 35 U.S.C. § 112, including by arguing that they are indefinite, not supported by the written description and/or not enabled. Accordingly, nothing stated herein shall be construed as a waiver of any argument available under 35 U.S.C. § 112. Nothing stated herein shall be construed as an admission or waiver with respect to the effective filing date of any of the Asserted Claims.


The following chart includes references to the structure, function, operation, and/or features of the technology described in ServiceGuard Manual. SAP reserve the right to present direct evidence of the structure, function, operation, and/or features of the technology described in ServiceGuard Manual (including product samples, source code, and/or witness testimony) and/or its prior art status, and to present alternative or additional documents concerning the structure, function, operation, and/or features of the technology described in the ServiceGuard Manual and/or its prior art status, as evidence that the technology described in ServiceGuard Manual anticipates and/or renders obvious the Asserted Claims of the ’182 Patent.

The cited portions of ServiceGuard Manual are merely illustrative, and Defendants reserve the right to rely on alternative or additional evidence, including uncited portions of ServiceGuard Manual. Where the chart below states that ServiceGuard Manual “describes” or “provides for” a limitation, such disclosure may be express, inherent, or implicit.

If not anticipated, one or more claims of the '182 Patent are rendered obvious in light of ServiceGuard Manual alone or ServiceGuard Manual in combination with the knowledge of a person of ordinary skill in the art.

In addition, if not anticipated, one or more claims of the '182 Patent are rendered obvious in light of the ServiceGuard Manual in combination with one or more of the references identified in SAP's Invalidity Contentions.

Claim Element	Claim Language	ServiceGuard Manual
1PRE	A data redundancy system, comprising:	ServiceGuard Manual discloses “a data redundancy system,” for example: ServiceGuard Manual at 22: MC/ServiceGuard allows you to create high availability clusters of HP 9000 Series 800 computers. A high availability computer system allows application services to continue in spite of a hardware or software failure. Highly available systems protect users from software failures as well as from failure of a system processing unit (SPU), disk, or local area network (LAN) component. In the event that one component fails, the redundant component takes over. MC/ServiceGuard and other high availability subsystems coordinate the transfer between components. An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption. ServiceGuard Manual at 22:


		Figure 1-1 Typical Cluster Configuration  Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
1(a)	a primary storage facility for storing a primary copy of data, the primary storage facility including;	ServiceGuard Manual discloses “a primary storage facility for storing a primary copy of data, the primary storage facility including,” for example: ServiceGuard Manual at 22: An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption.

ServiceGuard Manual at 23:

In the figure, node 1 (one of two SPU's) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package's applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line).

Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.


Figure 1-1 Typical Cluster Configuration



Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of

		ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
1(b)	a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and	ServiceGuard Manual discloses “a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility,” for example: ServiceGuard Manual at 22: An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption. ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package’s applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line). Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.

Figure 1-1 Typical Cluster Configuration




ServiceGuard Manual at 35:

Each node in a cluster has its own root disk, but each node is also physically connected to several other disks in such a way that more than one node can obtain access to the data and programs associated with a package it is configured for. This access is provided by a Storage Manager, such as Logical Volume Manager (LVM), VERITAS Volume Manager (VxVM), or VERITAS Cluster Volume Manager (CVM). A disk storage group can be activated by no more than one node at a time, but when the package is moved, the storage group can be activated by the adoptive node. All of the disks in the storage group owned by a package must be connected to the original node and to all possible adoptive nodes for that package. Disk storage is made redundant by using RAID or software mirroring.

Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.

1(c)	a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance;	ServiceGuard Manual discloses “a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance,” for example: ServiceGuard Manual at 22: An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption. ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package’s applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line). Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.
------	--	---


		Figure 1-1 Typical Cluster Configuration  Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
1(d)	a secondary storage facility for storing data that is redundant of the primary copy of the data, the secondary storage facility discrete from the primary storage facility and including:	ServiceGuard Manual discloses “a secondary storage facility for storing data that is redundant of the primary copy of the data, the secondary storage facility discrete from the primary storage facility and including,” for example: ServiceGuard Manual at 22: An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption.

ServiceGuard Manual at 23:

In the figure, node 1 (one of two SPU's) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package's applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line).


Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.

Figure 1-1 Typical Cluster Configuration



Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard

		Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
1(e)	a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility; and	ServiceGuard Manual discloses “a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility; and,” for example: ServiceGuard Manual at 22: An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption. ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package’s applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line). Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.


		Figure 1-1 Typical Cluster Configuration  Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
1(f)	a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance.	ServiceGuard Manual discloses “a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance,” for example: ServiceGuard Manual at 22: An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption.

ServiceGuard Manual at 23:

In the figure, node 1 (one of two SPU's) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package's applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line).

Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.

Figure 1-1 Typical Cluster Configuration




ServiceGuard Manual at 35:

Each node in a cluster has its own root disk, but each node is also physically connected to several other disks in such a way that more than one node can obtain access to the data and

		programs associated with a package it is configured for. This access is provided by a Storage Manager, such as Logical Volume Manager (LVM), VERITAS Volume Manager (VxVM), or VERITAS Cluster Volume Manager (CVM). A disk storage group can be activated by no more than one node at a time, but when the package is moved, the storage group can be activated by the adoptive node. All of the disks in the storage group owned by a package must be connected to the original node and to all possible adoptive nodes for that package. Disk storage is made redundant by using RAID or software mirroring. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	--

Claim Element	Claim Language	ServiceGuard Manual
2	The data redundancy system according to claim 1, wherein the secondary storage facility assumes the role of the primary storage facility in response to a failover event, the failover event including a communication traffic condition having a greater portion of request traffic originating closer to the secondary storage facility	ServiceGuard Manual discloses “The data redundancy system according to claim 1, wherein the secondary storage facility assumes the role of the primary storage facility in response to a failover event, the failover event including a communication traffic condition having a greater portion of request traffic originating closer to the secondary storage facility than to the primary storage facility,” for example: ServiceGuard Manual at 31: In the figure, a two-node MC/ServiceGuard cluster has one bridged net configured with both a primary and a standby LAN card for the data/heartbeat subnet (Subnet A). Another LAN card provides an optional dedicated heartbeat LAN. Note that the primary and standby LAN segments are connected by a hub to provide a redundant data/heartbeat subnet. Each node has its own IP address for this subnet. In case of a failure of a primary LAN card for the data/heartbeat subnet, MC/ServiceGuard will perform a local switch to the standby LAN card on the same node.

	than to the primary storage facility.	Redundant heartbeat is provided by the primary LAN and the dedicated LAN which are both carrying the heartbeat. In Figure 2-1, local switching is not needed for the dedicated heartbeat LAN, since there is already a redundant path via the other subnet. In case of data congestion on the primary LAN, the dedicated heartbeat LAN will prevent a false diagnosis of heartbeat failure. Each node has its own IP address for dedicated heartbeat LAN. You should verify that network traffic is not too high on the heartbeat/ data LAN. If traffic is too high, this LAN might not perform adequately in transmitting heartbeats if the dedicated heartbeat LAN fails. Figure 2-1 Redundant LANs  Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	--

Claim Element	Claim Language	ServiceGuard Manual
---------------	----------------	---------------------


3	The data redundancy system according to claim 1, wherein the second redundancy appliance assumes the role of the first redundancy appliance in response to a failover event selected from the group consisting of a manually initiated event and an automatically initiated event.	ServiceGuard Manual discloses “the data redundancy system according to claim 1, wherein the second redundancy appliance assumes the role of the first redundancy appliance in response to a failover event selected from the group consisting of a manually initiated event and an automatically initiated event,” for example: ServiceGuard Manual at 24: MC/ServiceGuard is designed to work in conjunction with other high availability products, such as MirrorDisk/UX or VERITAS Volume Manager, which provide disk redundancy to eliminate single points of failure in the disk subsystem; Event Monitoring Service (EMS), which lets you monitor and detect failures that are not directly handled by MC/ServiceGuard; disk arrays, which use various RAID levels for data protection; and HP-supported uninterruptible power supplies (UPS), such as HP PowerTrust, which eliminates failures related to power outage. These products are highly recommended along with MC/ServiceGuard to provide the greatest degree of availability. ServiceGuard Manual at 36: Disk mirroring is one method for providing data protection. The logical volumes used for MC/ServiceGuard packages should be mirrored. MC/ServiceGuard does not provide protection for data on your disks. This capability is provided for LVM storage with HP’s MirrorDisk/UX product, and for VxVM and CVM with the VERITAS Volume Manager (B9116AA). When you configure logical volumes using software mirroring, the members of each mirrored set contain exactly the same data. If one disk should fail, the storage manager will automatically keep the data available by accessing the other mirror. Three-way mirroring in LVM (or additional plexes with VxVM) may be used to allow for online backups or even to provide an additional level of high availability. To protect against FibreChannel or SCSI bus failures, each copy of the data must be accessed by a separate bus; that is, you cannot have all copies of the data on disk drives connected to the same bus. While it is valuable to mirror your root disks, it is critical for high availability that you mirror your data disks. In the event of damage to a root disk of one SPU, the other SPUs in the cluster will take over control of the applications. However, if a data disk is damaged, and not
---	---	---

		mirrored, any application which depends on that disk will not be available until the problem with the disk is repaired. Even then, it may not be possible to recover the data on that disk. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	--

Claim Element	Claim Language	ServiceGuard Manual
4	The data redundancy system according to claim 1, wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage by forwarding a write record to the second redundancy appliance for each write request, the second redundancy appliance sending acknowledgements to the first redundancy appliance in batches in response to receiving write records.	ServiceGuard Manual discloses “the data redundancy system according to claim 1, wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage by forwarding a write record to the second redundancy appliance for each write request, the second redundancy appliance sending acknowledgements to the first redundancy appliance in batches in response to receiving write records,” for example: ServiceGuard Manual at 337: As an alternative to using a TPM, queue up requests when the server is unavailable. Rather than notifying the user when a server is unavailable, the user request is queued up and transmitted later when the server becomes available again. Message queueing software ensures that messages of any kind, not necessarily just transactions, are delivered and acknowledged. Message queueing is useful only when the user does not need or expect response that the request has been completed (i.e, the application is not interactive). Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of

		ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	--

Claim Element	Claim Language	ServiceGuard Manual
5	The data redundancy system according to claim 4, wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance.	ServiceGuard Manual discloses “the data redundancy system according to claim 4, wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance,” for example: ServiceGuard Manual at 337: As an alternative to using a TPM, queue up requests when the server is unavailable. Rather than notifying the user when a server is unavailable, the user request is queued up and transmitted later when the server becomes available again. Message queueing software ensures that messages of any kind, not necessarily just transactions, are delivered and acknowledged. Message queueing is useful only when the user does not need or expect response that the request has been completed (i.e, the application is not interactive). Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
Claim Element	Claim Language	ServiceGuard Manual
6	The data redundancy system according to claim 1, wherein the first	ServiceGuard Manual discloses “the data redundancy system according to claim 1, wherein the first redundancy appliance is disjoint from the second redundancy appliance,” for example:

	redundancy appliance is disjoint from the second redundancy appliance.	ServiceGuard Manual at 22: MC/ServiceGuard allows you to create high availability clusters of HP 9000 Series 800 computers. A high availability computer system allows application services to continue in spite of a hardware or software failure. Highly available systems protect users from software failures as well as from failure of a system processing unit (SPU), disk, or local area network (LAN) component. In the event that one component fails, the redundant component takes over. MC/ServiceGuard and other high availability subsystems coordinate the transfer between components. An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption. ServiceGuard Manual at 22: Figure 1-1 Typical Cluster Configuration  The diagram illustrates a typical cluster configuration. It features two main nodes, 'node1' and 'node2', represented as large teal boxes. Each node contains an orange box labeled 'pkgA' and 'pkgB' respectively. To the left of node1 is a 'root mirror' disk, and to the right of node2 is a 'root mirror' disk. Between the two nodes are two sets of disks: 'pkgA disks' and 'pkgA mirrors' on the left, and 'pkgB disks' and 'pkgB mirrors' on the right. All these components are connected to a central horizontal line representing a network or bus. Below this line is a 'hub' box, which is connected to three client computers (represented as small blue boxes) at the bottom. The entire setup is connected to a common ground line at the very bottom.
--	---	---


		Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	--

Claim Element	Claim Language	ServiceGuard Manual
7	The data redundancy system according to claim 1, wherein the second redundancy appliance maintains a shadow copy of data stored by the first redundancy appliance wherein the second redundancy appliance accesses the shadow copy of the data when the second redundancy appliance assumes the role of the first redundancy appliance.	ServiceGuard Manual discloses “the data redundancy system according to claim 1, wherein the second redundancy appliance maintains a shadow copy of data stored by the first redundancy appliance wherein the second redundancy appliance accesses the shadow copy of the data when the second redundancy appliance assumes the role of the first redundancy appliance,” for example: ServiceGuard Manual at 22: An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption. ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package’s applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks

(dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line).

Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.

Figure 1-1 Typical Cluster Configuration




ServiceGuard Manual at 35:

Each node in a cluster has its own root disk, but each node is also physically connected to several other disks in such a way that more than one node can obtain access to the data and programs associated with a package it is configured for. This access is provided by a Storage Manager, such as Logical Volume Manager (LVM), VERITAS Volume Manager (VxVM), or VERITAS Cluster Volume Manager (CVM). A disk storage group can be activated by no more than one node at a time, but when the package is moved, the storage group can be activated by the adoptive node. All of the disks in the storage group owned by a package must be connected to the original node and to all possible adoptive nodes for that package. Disk storage is made redundant by using RAID or software mirroring.

		Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
Claim Element	Claim Language	ServiceGuard Manual
9	The data redundancy system according to claim 1, wherein the first redundancy appliance forwards redundant data to the secondary storage facility as send batches.	ServiceGuard Manual discloses “the data redundancy system according to claim 1, wherein the first redundancy appliance forwards redundant data to the secondary storage facility as send batches,” for example: ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package’s applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line). Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.

Figure 1-1 Typical Cluster Configuration



ServiceGuard Manual at 344-345:

2. Be sure to test all combinations of application load during the testing. Repeat the failover processes under different application states such as heavy user load versus no user load, batch jobs vs online transactions, etc.

Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.

Claim Element	Claim Language	ServiceGuard Manual
10	The data redundancy system according to claim 9, wherein a send batch size is determined	ServiceGuard Manual discloses “the data redundancy system according to claim 9, wherein a send batch size is determined according to either a time interval for collecting transactions into the send batch or a predetermined quantity of transactions to be included in the send batch,” for example:

	according to either a time interval for collecting transactions into the send batch or a predetermined quantity of transactions to be included in the send batch.	ServiceGuard Manual at 96: There are tradeoffs in only using a single subnet. If the heartbeat is not heard for more than a NODE_TIMEOUT interval, the cluster will re-form. Depending on the length of the disruption, this could cause one or all of the cluster nodes to TOC. This is most commonly seen when there is heavy traffic on the subnet preventing the heartbeat packets from being transmitted. This can also be seen when there are network configuration problems or network hardware problems that are not easily recovered and may not be detected by ServiceGuard or APA. The most common way to avoid problems in these scenarios is to utilize multiple heartbeat networks. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	--


Claim Element	Claim Language	ServiceGuard Manual
11	The data redundancy system according to claim 9, wherein a send batch size is determined adaptively based on available communication bandwidth between the primary storage facility and the secondary storage facility.	ServiceGuard Manual discloses “the data redundancy system according to claim 9, wherein a send batch size is determined adaptively based on available communication bandwidth between the primary storage facility and the secondary storage facility,” for example: ServiceGuard Manual at 69: You can specify a registered resource for a package by selecting it from the list of available resources displayed in the SAM package configuration area. The size of the list displayed by SAM depends on which resource monitors have been registered on your system. Alternatively, you can obtain information about registered resources on your system by using the command /opt/resmon/bin/resls. For additional information, refer to the man page for resls(1m).

		ServiceGuard Manual at 325-326: Some databases permit logs to be buffered in memory to increase online performance. Of course, when a failure occurs, any in-flight transaction will be lost. However, minimizing the size of this in-memory log will reduce the amount of completed transaction data that would be lost in case of failure. Keeping the size of the on-disk log small allows the log to be archived or replicated more frequently, reducing the risk of data loss if a disaster were to occur. There is, of course, a trade-off between online performance and the size of the log. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	--

Claim Element	Claim Language	ServiceGuard Manual
12	The data redundancy system according to claim 1, wherein a first write transaction in a send batch at the primary storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction.	ServiceGuard Manual discloses “the data redundancy system according to claim 1, wherein a first write transaction in a send batch at the primary storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction,” for example: ServiceGuard Manual at 95: CVM allows you to activate storage on one node at a time, or you can perform write activation on one node and read activation on another node at the same time (for example, allowing backups). CVM provides full mirroring and dynamic multipathing (DMP) for clusters. CVM can be used in clusters that: • will run applications that require fast disk group activation after package failover. • require activation on more than one node at a time, for example to perform a backup from one node while a package using the volume is active on another node. In this case, the

		package using the disk group would have the disk group active in exclusive write mode while the node that is doing the backup would have the disk group active in shared read mode. • are configured with only a single heartbeat subnet. CVM is supported on 4 nodes or fewer at this release. Shared storage devices must be connected to all nodes in the cluster, whether or not the node accesses data on the device. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
Claim Element	Claim Language	ServiceGuard Manual
13	The data redundancy system according to claim 1, wherein the third redundancy appliance stores redundant data to the storage associated with the secondary facility according to receive batches.	ServiceGuard Manual discloses “the data redundancy system according to claim 1, wherein the third redundancy appliance stores redundant data to the storage associated with the secondary facility according to receive batches,” for example: ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package’s applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line). Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.

Figure 1-1 Typical Cluster Configuration




Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.

Claim Element	Claim Language	ServiceGuard Manual
14	The data redundancy system according to claim 13, wherein the third redundancy appliance stores a receive batch after the receive batch is determined safe to apply.	ServiceGuard Manual discloses “the data redundancy system according to claim 13, wherein the third redundancy appliance stores a receive batch after the receive batch is determined safe to apply,” for example: ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package’s applications, and a mirror copy of the data. Note that both nodes are physically

connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line).

Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.

Figure 1-1 Typical Cluster Configuration



Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.

Claim Element	Claim Language	ServiceGuard Manual
15	The data redundancy system according to claim 14, wherein a receive batch is determined safe to apply when a data integrity test is successfully performed.	ServiceGuard Manual discloses “the data redundancy system according to claim 14, wherein a receive batch is determined safe to apply when a data integrity test is successfully performed,” for example: ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by the package’s applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line). Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.

		Figure 1-1 Typical Cluster Configuration Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	--


Claim Element	Claim Language	ServiceGuard Manual
16	The data redundancy system according to claim 13, wherein a first write transaction in a first receive batch at the secondary storage facility is replaced with a second write transaction in a second receive batch	ServiceGuard Manual discloses “the data redundancy system according to claim 13, wherein a first write transaction in a first receive batch at the secondary storage facility is replaced with a second write transaction in a second receive batch wherein the second write transaction overwrites data written by the first write transaction and the first and second receive batches are combined,” for example: ServiceGuard Manual at 23: In the figure, node 1 (one of two SPU’s) is running package A, and node 2 is running package B. Each package has a separate group of disks associated with it, containing data needed by

wherein the second write transaction overwrites data written by the first write transaction and the first and second receive batches are combined.

the package's applications, and a mirror copy of the data. Note that both nodes are physically connected to both groups of mirrored disks. However, only one node at a time may access the data for a given group of disks. In the figure, node 1 is shown with exclusive access to the top two disks (solid line), and node 2 is shown as connected without access to the top disks (dotted line). Similarly, node 2 is shown with exclusive access to the bottom two disks (solid line), and node 1 is shown as connected without access to the bottom disks (dotted line).


Mirror copies of data provide redundancy in case of disk failures. In addition, a total of four data buses are shown for the disks that are connected to node 1 and node 2. This configuration provides the maximum redundancy and also gives optimal I/O performance, since each package is using different buses.


Figure 1-1 Typical Cluster Configuration




Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.


Claim Element	Claim Language	ServiceGuard Manual
17PRE	A data redundancy method for storing a primary copy of data at a primary storage facility and for storing data that is redundant of the primary copy at a secondary storage facility, the method comprising:	ServiceGuard Manual discloses “A data redundancy method for storing a primary copy of data at a primary storage facility and for storing data that is redundant of the primary copy at a secondary storage facility, the method comprising,” for example: ServiceGuard Manual at 22: MC/ServiceGuard allows you to create high availability clusters of HP 9000 Series 800 computers. A high availability computer system allows application services to continue in spite of a hardware or software failure. Highly available systems protect users from software failures as well as from failure of a system processing unit (SPU), disk, or local area network (LAN) component. In the event that one component fails, the redundant component takes over. MC/ServiceGuard and other high availability subsystems coordinate the transfer between components. An MC/ServiceGuard cluster is a networked grouping of HP 9000 series 800 servers (host systems known as nodes) having sufficient redundancy of software and hardware that a single point of failure will not significantly disrupt service. Application services (individual HP-UX processes) are grouped together in packages; in the event of a single service, node, network, or other resource failure, MC/ServiceGuard can automatically transfer control of the package to another node within the cluster, allowing services to remain available with minimal interruption. ServiceGuard Manual at 22:

		Figure 1-1 Typical Cluster Configuration  Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
17(a)	receiving a sequence of write requests at a first redundancy appliance at the primary storage facility;	ServiceGuard Manual discloses “receiving a sequence of write requests at a first redundancy appliance at the primary storage facility,” for example: ServiceGuard Manual at 23-24: Under normal conditions, a fully operating MC/ServiceGuard cluster simply monitors the health of the cluster's components while the packages are running on individual nodes. Any host system running in the MC/ServiceGuard cluster is called an active node. When you create the package, you specify a primary node and one or more adoptive nodes. When a node or its network communications fails, MC/ServiceGuard can transfer control of the package to the next available adoptive node. This situation is shown in Figure 1-2.


		After this transfer, the package typically remains on the adoptive node as long the adoptive node continues running. If you wish, however, you can configure the package to return to its primary node as soon as the primary node comes back online. Alternatively, you may manually transfer control of the package back to the primary node at the appropriate time. ServiceGuard Manual at 24: Figure 1-2 Typical Cluster After Failover  Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
17(b)	storing data for the sequence of write requests in primary storage at the primary storage facility;	ServiceGuard Manual discloses “storing data for the sequence of write requests in primary storage at the primary storage facility,” for example: ServiceGuard Manual at 23-24: Under normal conditions, a fully operating MC/ServiceGuard cluster simply monitors the health of the cluster's components while the packages are running on individual nodes. Any host system running in the MC/ServiceGuard cluster is called an active node. When you create the package, you specify a primary node and one or more adoptive nodes. When a node

		or its network communications fails, MC/ServiceGuard can transfer control of the package to the next available adoptive node. This situation is shown in Figure 1-2. After this transfer, the package typically remains on the adoptive node as long the adoptive node continues running. If you wish, however, you can configure the package to return to its primary node as soon as the primary node comes back online. Alternatively, you may manually transfer control of the package back to the primary node at the appropriate time. ServiceGuard Manual at 24: Figure 1-2 Typical Cluster After Failover  Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
17(c)	logging at least a sequence of write requests for transfer to the secondary storage facility;	ServiceGuard Manual discloses “logging at least a sequence of write requests for transfer to the secondary storage facility,” for example: ServiceGuard Manual at 49:

		Each of these daemons logs to the /var/adm/syslog/syslog.log file except for the quorum server, which logs to the standard output (it is suggested you redirect output to a file named /var/adm/qs/qs.log) and /usr/sbin/cmomd, which logs to /var/opt/cmom/cmomd.log. ServiceGuard Manual at 79: At any step along the way, an error will result in the script exiting abnormally (with an exit code of 1). Also, if the halt script execution is not complete before the time specified in the HALT_SCRIPT_TIMEOUT, the package manager will kill the script. During halt script execution, messages are written to a log file in the same directory as the halt script. This log has the same name as the halt script and the extension .log. Normal starts are recorded in the log, together with error messages or warnings related to halting the package. ServiceGuard Manual at 177: You need to initialize the physical disks that will be employed in VxVM disk groups. To initialize a disk, log on to one node in the cluster, then use the vxdiskadm program to initialize multiple disks, or use the vxdisksetup command to initialize one disk at a time, as in the following example: # /usr/lib/vxvm/bin/vxdisksetup -i c0t3d2 ServiceGuard Manual at 194: You need to initialize the physical disks that will be employed in CVM disk groups. If a physical disk has been previously used with LVM, you should use the pvremove command to delete the LVM header data from all the disks in the volume group (this is not necessary if you have not previously used the disk with LVM). To initialize a disk for CVM, log on to the master node, then use the vxdiskadm program to initialize multiple disks, or use the vxdisksetup command to initialize one disk at a time, as in the following example: # /usr/lib/vxvm/bin/vxdisksetup -i /dev/dsk/c0t3d2 Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard
--	--	--

		Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
17(d)	forwarding the logged write requests as send batches to a second redundancy appliance at the secondary storage facility disjoint from the primary storage facility, the second storage facility batching the received write request as receive batches;	ServiceGuard Manual discloses “forwarding the logged write requests as send batches to a second redundancy appliance at the secondary storage facility disjoint from the primary storage facility, the second storage facility batching the received write request as receive batches,” for example: ServiceGuard Manual at 23-24: Under normal conditions, a fully operating MC/ServiceGuard cluster simply monitors the health of the cluster's components while the packages are running on individual nodes. Any host system running in the MC/ServiceGuard cluster is called an active node. When you create the package, you specify a primary node and one or more adoptive nodes. When a node or its network communications fails, MC/ServiceGuard can transfer control of the package to the next available adoptive node. This situation is shown in Figure 1-2. After this transfer, the package typically remains on the adoptive node as long the adoptive node continues running. If you wish, however, you can configure the package to return to its primary node as soon as the primary node comes back online. Alternatively, you may manually transfer control of the package back to the primary node at the appropriate time. ServiceGuard Manual at 24: Figure 1-2 Typical Cluster After Failover 

		Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
17(e)	storing data associated with the receive batches in secondary storage at the secondary storage facility;	ServiceGuard Manual discloses “storing data associated with the receive batches in secondary storage at the secondary storage facility,” for example: ServiceGuard Manual at 23-24: Under normal conditions, a fully operating MC/ServiceGuard cluster simply monitors the health of the cluster's components while the packages are running on individual nodes. Any host system running in the MC/ServiceGuard cluster is called an active node. When you create the package, you specify a primary node and one or more adoptive nodes. When a node or its network communications fails, MC/ServiceGuard can transfer control of the package to the next available adoptive node. This situation is shown in Figure 1-2. After this transfer, the package typically remains on the adoptive node as long the adoptive node continues running. If you wish, however, you can configure the package to return to its primary node as soon as the primary node comes back online. Alternatively, you may manually transfer control of the package back to the primary node at the appropriate time. ServiceGuard Manual at 24:

		Figure 1-2 Typical Cluster After Failover  Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
17(f)	replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction; and	ServiceGuard Manual discloses “replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction,” for example: ServiceGuard Manual at 95: CVM allows you to activate storage on one node at a time, or you can perform write activation on one node and read activation on another node at the same time (for example, allowing backups). CVM provides full mirroring and dynamic multipathing (DMP) for clusters. CVM can be used in clusters that: • will run applications that require fast disk group activation after package failover. • require activation on more than one node at a time, for example to perform a backup from one node while a package using the volume is active on another node. In this case, the package using the disk group would have the disk group active in exclusive write mode while the node that is doing the backup would have the disk group active in shared read mode. • are configured with only a single heartbeat subnet.

		CVM is supported on 4 nodes or fewer at this release. Shared storage devices must be connected to all nodes in the cluster, whether or not the node accesses data on the device. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
17(g)	the second redundancy appliance acting as the first redundancy appliance in response to a failover event.	ServiceGuard Manual discloses “the second redundancy appliance acting as the first redundancy appliance in response to a failover event,” for example: ServiceGuard Manual at 233: Failover and Failback Policies Packages can be configured with one of two values for the FAILOVER_POLICY parameter: • CONFIGURED_NODE. The package fails over to the next node in the node list in the package configuration file. • MIN_PACKAGE_NODE. The package fails over to the node in the cluster with the fewest running packages on it. Packages can also be configured with one of two values for the FAILBACK_POLICY parameter: • AUTOMATIC. With this setting, a package, following a failover, returns to its primary node when the primary node becomes available again. • MANUAL. With this setting, a package, following a failover, must be moved back to its original node by a system administrator. Failover and failback policies are displayed in the output of the cmviewcl -v command. ServiceGuard Manual at 334:

		Use separate volume groups for each application that uses data. If the application doesn't use disk, it is not necessary to assign it a separate volume group. A volume group (group of disks) is the unit of storage that can move between nodes. The greatest flexibility for load balancing exists when each application is confined to its own volume group, i.e., two applications do not share the same set of disk drives. If two applications do use the same volume group to store their data, then the applications must move together. If the applications' data stores are in separate volume groups, they can switch to different nodes in the event of a failover. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
--	--	---

Claim Element	Claim Language	ServiceGuard Manual
18PRE	A data redundancy system, comprising:	ServiceGuard Manual discloses “a data redundancy system, comprising,” for example: <i>See analysis of 1[pre]</i> Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
18(a)	a primary storage facility including at least one primary redundancy appliance for receiving a	ServiceGuard Manual discloses “a primary storage facility including at least one primary redundancy appliance for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility,” for example:

	sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility, the primary storage facility batching the write requests as send batches for transfer to a secondary storage facility; and	<i>See</i> analysis of limitations 1[b], 1[c], 9. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.
18(b)	the secondary storage facility disjoint from the primary storage facility including at least one secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in storage associated with the secondary storage facility, the second storage facility batching the received write request as receive batches, the second redundancy appliance acting as the primary redundancy	ServiceGuard Manual discloses “the secondary storage facility disjoint from the primary storage facility including at least one secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in storage associated with the secondary storage facility,” for example: <i>See</i> analysis of limitations 17[d], 17[g]. Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.

	appliance in response to a failover event;	
18(c)	wherein a first write transaction in a receive batch at the second storage facility is replaced with a second write transaction where the second write transaction over writes data written by the first write transaction.	ServiceGuard Manual discloses “wherein a first write transaction in a receive batch at the second storage facility is replaced with a second write transaction where the second write transaction over writes data written by the first write transaction,” for example: <i>See analysis of limitation 17[f]</i> Viewing ServiceGuard Manual as a whole and view of contemporary prior art, this element would have been inherent or obvious to one of ordinary skill in the art. To the extent it is found that the element is not expressly disclosed in ServiceGuard Manual, ServiceGuard Manual renders it obvious, either alone, in combination with the knowledge of a person of ordinary skill in the art, and/or in combination with other prior art references identified in the cover pleading or herein.