

**UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
WACO DIVISION**

INTELLECTUAL VENTURES II LLC,)	
)	Civil Action No. 6:21-cv-1298
Plaintiffs,)	
)	
v.)	
)	
HEWLETT PACKARD ENTERPRISE CO.,))	
)	
Defendant.)	JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff, Intellectual Ventures II LLC (“IV”), for its complaint against defendant, Hewlett Packard Enterprise Company (“HPE”), hereby alleges as follows:

THE PARTIES

1. Intellectual Ventures II LLC (“Intellectual Ventures II” or “IV”) is a Delaware limited liability company having its principal place of business located at 3150 139th Avenue SE, Bellevue, Washington 98005.
2. Upon information and belief, HPE is a Delaware corporation with its principal executive offices located at 11445 Compaq Center West Drive, Houston, Texas 77070.
3. Upon information and belief, HPE has regular and established places of business in this District, including a fifty-two (52) acre campus at 14321 Tandem Boulevard, Austin, Texas, and a lease for another 27,326 square foot office at Paloma Ridge, 13620 FM 620 Austin, Texas 78717. HPE also has an office at 6080 Tennyson Parkway, Suite 400, Plano, Texas 75024.
4. Upon information and belief, HPE’s global headquarters is located in Houston, Texas.

5. Upon information and belief, HPE may be served with process through its registered agent, CT Corporation System, 1999 Bryan Street, Suite 900, Dallas, Texas 75201.

6. Upon information and belief, the accused GreenLake products are headquartered in Texas.

JURISDICTION

7. IV brings this action for patent infringement pursuant to 35 U.S.C. § 271, *et seq.* This Court has subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a).

8. This Court has general jurisdiction over HPE because HPE is engaged in substantial and not isolated activity at its regular and established places of business within this judicial district. This Court has specific jurisdiction over HPE because HPE has committed acts of infringement giving rise to this action within this judicial district and has established more than minimum contacts within this judicial district, such that the exercise of jurisdiction over HPE in this Court would not offend traditional notions of fair play and substantial justice.

9. Venue is proper in this judicial district pursuant to 28 U.S.C. §§ 1391(b)-(c) and 1400(b) because HPE maintains regular and established places of business and has committed acts of patent infringement within this judicial district.

FACTUAL BACKGROUND

10. Intellectual Ventures Management, LLC (“Intellectual Ventures”) was founded in 2000. Since then, Intellectual Ventures has been involved in the invention business. Intellectual Ventures fosters inventions and facilitates the filing of patent applications for those inventions; collaborates with others to develop and patent inventions; and acquires and licenses patents from individual inventors, universities, corporations, and other institutions. A significant aspect of Intellectual Ventures’ business is managing the plaintiff in this case, Intellectual Ventures II.

11. One founder of Intellectual Ventures is Nathan Myhrvold, who worked at Microsoft from 1986 until 2000 in a variety of executive positions, culminating in his appointment as the company's first Chief Technology Officer (“CTO”) in 1996. While at Microsoft, Dr. Myhrvold founded Microsoft Research in 1991 and was one of the world’s foremost software experts. Between 1986 and 2000, Microsoft became the world’s largest technology company.

12. Under Dr. Myhrvold’s leadership, IV acquired more than 70,000 patents covering many important inventions of the Internet era. Many of these inventions coincided with Dr. Myhrvold’s successful tenure at Microsoft.

13. One area of particular importance in today’s Big Data driven technology sector is the emergence of cluster computing, such as, *e.g.*, high-performance computing (“HPC”), high availability (“HA”), and load balancing (“LB”), as a pay-as-you-go service. Traditionally, cluster computing—including HPC, HA, and LB—was difficult to implement, costly to maintain, and inflexible in its configuration. For example, the various implementations of computing clusters could be targeted at only one configuration—*e.g.*, HPC or HA, but not both—and could not be reconfigured from one configuration to the other once implemented. These barriers prevented wide-spread adoption of the technology and relegated it to huge labs operated by multi-national organizations, wealthy governments, or very large companies and research institutions. While the proliferation of the public cloud over the past decade revolutionized how individuals and organizations consume data and fostered a huge migration of certain applications, data, and services off-premises into the cloud, it did little to overcome the limitations of cluster computing in general and HPC implementations in particular. The public cloud was not sufficient for many HPC, HA, or LB applications, which often require specific configurations tailored to a particular

task, like high availability, at one point in time but might require another task, like HPC, either simultaneously or at a later date.

14. To overcome these traditional implementation and affordability obstacles there has been a large push to provide HPC (and other cluster computing implementations) as a service, removing the burden of purchasing and maintaining a data center with static configurations of compute nodes from the shoulders of those who desire an HPC environment. Service providers began investing in the infrastructure required for such massive computing, incorporating advanced hardware and software—often relying heavily on increasing bandwidth and virtualization—and then offering customizable cluster implementations to customers that would have found the complexity and expenditure of traditional systems insurmountable.

15. HPE makes, uses, and sells cluster computing and HPC platforms, including hardware and software, with a particular focus on providing such platforms as a service. One of the fastest growing implementations of HPE's cluster computing offerings is HPE GreenLake for HPC, which provides hosted and managed HPC environments to consumers as a pay-as-you-go service. This service utilizes HPE's specially built Apollo and ProLiant hardware and includes purpose-built software like the HPE Performance Cluster Manager, HPE container platform, and high-performance storage and networking capabilities.

16. Upon information and belief, HPE customers of one or more of the Accused Products include the National Security Agency, and the University of Ghent.

17. Upon information and belief, HPE customers of one or more of the Accused Products also include University of Houston, Cullen College of Engineering.

18. Upon information and belief, ScaleMatrix's Dallas Data Center provides one or more of the Accused Products as a service.

THE PATENT-IN-SUIT

19. On October 26, 2010, the PTO issued United States Patent No. 7,822,841 (“the ’841 patent”), titled METHOD AND SYSTEM FOR HOSTING MULTIPLE, CUSTOMIZED COMPUTING CLUSTERS. The ’841 patent is valid and enforceable. A copy of the ’841 patent is attached as Exhibit A.

20. Intellectual Ventures II LLC is the owner and assignee of all rights, title, and interest in and to the ’841 patent, and holds all substantial rights therein, including the rights to grant licenses, to exclude others, and to enforce and recover past damages for infringement of that patent.

21. The inventions claimed in the ’841 patent were conceived by Jeffrey B. Franklin, a serial entrepreneur and Co-Founder of Modern Grids, Inc., an early pioneer in custom hosted high-performance cluster computing. Mr. Franklin received his bachelor’s degree in Computer Science and Computer Engineering from Kettering University and subsequently worked for industry leaders, such as Rockwell Automation and Ansys, Inc., before co-founding Modern Grids, Inc. and joining his current company, Limitless Computing, as CTO. Mr. Franklin is a named inventor on numerous patents over his 25-year career in the technology industry and, in addition to designing and developing pioneering cloud technology and foundational frameworks for the same from the ground up, has played a key role in bringing over ten AR/VR/Cloud products and services to market, including some used by industry leaders like Apple and Google. He currently resides and works in Boulder, Colorado.

22. The ’841 patent is directed to systems, methods, and/or apparatus for hosting computing clusters for clients, each cluster including a custom or differing set of configurations

and including a monitoring system that monitors the clusters for operational problems on a per-cluster and per-node basis and controls client access to the clusters.

COUNT I

(HPE's Infringement of U.S. Patent No. 7,822,841)

23. The preceding paragraphs are reincorporated by reference as if fully set forth herein.

24. The inventions claimed in the '841 patent, taken alone or in combination, were not well-understood, routine, or conventional to one of ordinary skill in the art at the time of the invention. Rather, the '841 patent claims and teaches, *inter alia*, an improved way to provide high performance cluster computing for a wide variety of uses across multiple industry verticals. The inventions improved upon then-existing HPC and cluster computing technology by developing a hosted computing infrastructure that is dynamically customizable and highly scalable, while also maintaining the robust security required for many HPC applications.

25. The inventions claimed in the '841 patent represent technical solutions to an unsolved technological problem. The written description of the '841 patent describes, in technical detail, each of the limitations in the claims, allowing a person of skill in the art to understand what those limitations cover, and therefore what was claimed, and also understand how the non-conventional and non-generic ordered combination of the elements of the claims differ markedly from what had been performed in the industry prior to the inventions of the '841 patent. More specifically, the claims of the '841 patent recite a computer system for hosting computing clusters comprising a private communications network linked to a public communications network with a first cluster comprising computing resources, including a processor, in a first configuration, where the cluster is linked to the private network. Furthermore, the system includes a second cluster also comprising a set of computing resources, including a processor, in a second configuration, the

second cluster also linked to the private network. The claim specifies that the first and second cluster configurations differ in that the first cluster provides a first computing environment for performing a first task and the second cluster provides a second computing environment for performing a second task. Additionally, the system includes a monitoring system that monitors the operations of the first and second clusters that identifies operational and connectivity problems and issues corresponding alerts indicating which of the clusters the alerts apply to. The monitoring system further comprises monitors for each node in the first and second clusters checking for hardware and software problems in the node and reporting identified problems to the main monitor.

26. The system covered by the asserted claims, therefore, differs markedly from the prior systems in use at the time of this invention, which, *inter alia*, lacked the claimed combination of two different, common-network-connected cluster configurations being monitored by a main monitor and a node-level monitor for identifying both cluster level operational and connectivity problems as well as node level hardware and software problems.

27. The '841 patent is drawn to solving a specific, technical problem arising in the context of hosted and customizable cluster computing systems. Consistent with the problem addressed being rooted in such HPC technology, the solutions disclosed in the '841 patent consequently are also rooted in that same technology and cannot be performed with pen and paper or in the human mind.

28. HPE has directly infringed, and continues to directly infringe, literally and/or by the doctrine of equivalents, individually and/or jointly, at least claim 1 of the '841 patent by making, using, testing, selling, offering for sale, and/or importing into the United States products and/or services covered by one or more claims of the '841 patent. HPE's products and/or services that infringe the '841 patent include, but are not limited to, the HPE GreenLake for High

Performance Computing (HPC), HPE GreenLake for Containers, GreenLake for Private Cloud, HPC Performance Cluster Manager, and related HPE Apollo and ProLiant Systems, as well as any other HPE products and/or services, either alone or in combination, that operate in substantially the same manner (together the “Accused ’841 Products” or “Accused Products”).

29. Claim 1 of the ’841 patent is reproduced below:

*1. A computer system for hosting computing clusters for clients, comprising:
a private communications network linked to a public communications network;*

a first cluster comprising a set of computing resources, including at least one hardware processor, in a first configuration, wherein the first cluster is communicatively linked to the private communications network;

a second cluster comprising a set of computing resources, including at least one hardware processor, in a second configuration, wherein the second cluster is communicatively linked to the private communications network;

a monitoring system monitoring operations of the first and second clusters, identifying operational and connectivity problems, and issuing an alert in response to the identified problems indicating a corresponding one of the first and second clusters associated with the identified problems;

wherein the first configuration differs from the second configuration and wherein the first configuration provides a first computing environment for performing a first client task and the second configuration provides a second computing environment for performing a second client task;

wherein the monitoring system comprises a main monitor that operates to monitor the first and second clusters to identify the operation and connectivity problems and further comprises monitors for each node of the first and second clusters operating to check for hardware and software problems within a particular node and to report the hardware and software problems to the main monitor.

30. The Accused ’841 Products provide a computer system for hosting computing clusters for clients. As one non-limiting example, the Accused ’841 Products include the HPE

HPC platform, including but not limited to HPE GreenLake for HPC, which includes a Performance Cluster Manager, as seen below:

	Small configuration	Medium configuration	Large configuration
Who is this recommended for?	<ul style="list-style-type: none"> Number of racks: 1 Total physical cores: 1,280 Compute: 20 HPE Apollo 2000 XL225n Q CPUs, 64 cores, 512 GB Scratch storage: Local NVME Shared storage: 8 HPE DL325 storage nodes with 232 TB raw capacity Network: High-performance 100 GbE Aruba GPU (Optional): NVIDIA A100 GPU with HPE Apollo 6500 servers 	<ul style="list-style-type: none"> Number of racks: 1 Total physical cores: 4,352 Compute: 68 HPE Apollo 2000 XL225n Q CPUs, 64 cores, 512 GB Scratch storage: Local NVME Shared storage: 16 HPE DL325 storage nodes with 464 TB raw capacity Network: High-performance 100 GbE Aruba GPU (Optional): NVIDIA A100 GPU with HPE Apollo 6500 servers 	<ul style="list-style-type: none"> Number of racks: 8 Total physical cores: 8,560 Compute: 140 HPE Apollo 2000 XL225n Q CPUs, 64 cores, 512 GB Scratch storage: Local NVME Shared storage: 32 HPE DL325 storage nodes with 812 TB raw capacity Network: High-performance 100 GbE Aruba GPU (Optional): NVIDIA A100 GPU with HPE Apollo 6500 servers
Software stack	<ul style="list-style-type: none"> Singularity Open Source (container runtime) Singularity PRO (optional) Slurm SLES 15 SP2 (HPC-ready golden image) HPE Ezmeral Data Fabric (shared NFS) HPE Performance Cluster Manager 	<ul style="list-style-type: none"> Singularity Open Source (container runtime) Singularity PRO (optional) Slurm SLES 15 SP2 (HPC-ready golden image) HPE Ezmeral Data Fabric (shared NFS) HPE Performance Cluster Manager 	<ul style="list-style-type: none"> Singularity Open Source (container runtime) Singularity PRO (optional) Slurm SLES 15 SP2 (HPC-ready golden image) HPE Ezmeral Data Fabric (shared NFS) HPE Performance Cluster Manager
Deployment options	<ul style="list-style-type: none"> Customer data center Colocation facility 	<ul style="list-style-type: none"> Customer data center Colocation facility 	<ul style="list-style-type: none"> Customer data center Colocation facility
Control plane	HPE GreenLake Central orchestration with HPC-specific module to monitor and manage the HPC environment and jobs	HPE GreenLake Central orchestration with HPC-specific module to monitor and manage the HPC environment and jobs	HPE GreenLake Central orchestration with HPC-specific module to monitor and manage the HPC environment and jobs
What is metered	Usage is metered based on the compute (per hour) and storage (per GB) used by the nodes in a cluster. There are two Q2 meters used to calculate usage over the reserved capacity. <ul style="list-style-type: none"> Compute nodes on/off Raw capacity used on shared storage 	Usage is metered based on the compute (per hour) and storage (per GB) used by the nodes in a cluster. There are two Q2 meters used to calculate usage over the reserved capacity. <ul style="list-style-type: none"> Compute nodes on/off Raw capacity used on shared storage 	Usage is metered based on the compute (per hour) and storage (per GB) used by the nodes in a cluster. There are two Q2 meters used to calculate usage over the reserved capacity. <ul style="list-style-type: none"> Compute nodes on/off Raw capacity used on shared storage

New collaboration will tackle growing AI and data needs with HPE's industry-leading high performance computing solutions, delivered as a service

HOUSTON – September 1, 2021 – Hewlett Packard Enterprise (NYSE: HPE) today announced that it has been awarded a \$2B contract, that will be leveraged over a 10 year period, with the National Security Agency (NSA) to deliver HPE's high performance computing (HPC) technology as a service through the HPE GreenLake platform.

The new collaboration will enable the NSA to harness rapidly growing AI and data needs more efficiently to create insights and other forecasting and analysis with optimal performance. By using HPE's HPC solutions through the HPE GreenLake platform, which provides fully managed, secure cloud services on-premises, the NSA will benefit from an agile, flexible, and secure platform to meet their growing data management requirements.

Harnessing data growth with purpose-built HPC and AI solutions

The new service includes a combination of HPE Apollo systems and HPE ProLiant servers, which ingest and process high volumes of data, and support deep learning and artificial intelligence capabilities. As part of the HPE GreenLake service, HPE will build and manage the complete solution that will be hosted at a QTS data center, a hosting facility that delivers secure, compliant data center infrastructure and robust connectivity to support scaling of operations.

- **HPE GreenLake cloud services for bare metal:** Allows customers to discover, provision and manage compute and storage resources in HPE GreenLake Central, to run workloads directly on the server, or use VMs or containers in the mix that best fits their business. This provides customers with control, agility and a lower cost approach to build platforms tailored to their performance, cost, and configuration requirements.
- **HPE GreenLake cloud services for containers:** Using the same scalable building blocks as bare metal and virtual machines, the HPE GreenLake cloud service for containers offers solutions based on small, medium, large and extra-large capacity sizes for the HPE Ezmeral Container platform.

The cluster manager works with clusters that include the following components:

- An admin node. All clusters include an admin node. On HPE Apollo 9000 systems and HPE SGI 8600 systems, the admin node is an HPE ProLiant DL360 system.
- Compute nodes. Depending on the type of cluster, it can include the following types of compute nodes:
 - ICE compute nodes. These compute nodes, typically diskless, are under the control of an ICE leader node. Only the HPE SGI 8600 clusters and the SGI ICE series of clusters have ICE compute nodes.
 - Non-ICE compute nodes. These are compute nodes that are not under the control of an ICE leader node. They can be configured to be under the control of a scalable unit (SU) leader node, or they can be configured to be under the direct control of an admin node.

Clusters with ICE leader nodes often have non-ICE compute nodes deployed with user services. For example, non-ICE compute nodes can be configured as login nodes or gateway nodes.

Clusters can also be configured using the admin node and several non-ICE compute nodes. These clusters have no leader nodes.

- Leader nodes. Depending on the number of compute nodes in the clusters, HPE might recommend leader nodes. A cluster can have either ICE leader nodes or SU leader nodes.

31. Furthermore, the Accused '841 Products provide a private communications network linked to a public communications network. For example, the Accused Products include a private cluster management/internal network. Also, for example, the Accused Products provide connectivity links to public communications networks, such as the Internet, as seen below:

HPE Apollo clusters can have SU leader nodes. On a cluster with SU leader nodes, the admin node supports SU leader nodes in an **SU leader pool**. In the SU leader pool, a single SU leader node is configured with two other SU leader nodes into a computing trio. This trio of SU leader nodes forms a resilient unit that administers computing functions. You are required to configure SU leader nodes in multiples of three. For example, you can configure 3, 6, 9, or any number of SU leader nodes that is a multiple of three.




Figure 3: Cluster with SU leader nodes

leader node

A type of node that manages one or more ICE compute nodes or non-ICE compute nodes. The admin node and all the leader nodes attach to the management network directly.

non-ICE compute node

A type of compute node. The admin node and the non-ICE compute nodes attach to the management network directly.

The following figure shows clusters with leader nodes and without leader nodes.



In a cluster without leader nodes, the admin node and the compute nodes attach directly to the management network. This cluster includes the admin node and several non-ICE compute nodes. All nodes attach directly to the management network, and all compute nodes can communicate with the admin node. Several HPE Apollo models are configured into clusters without leader nodes.

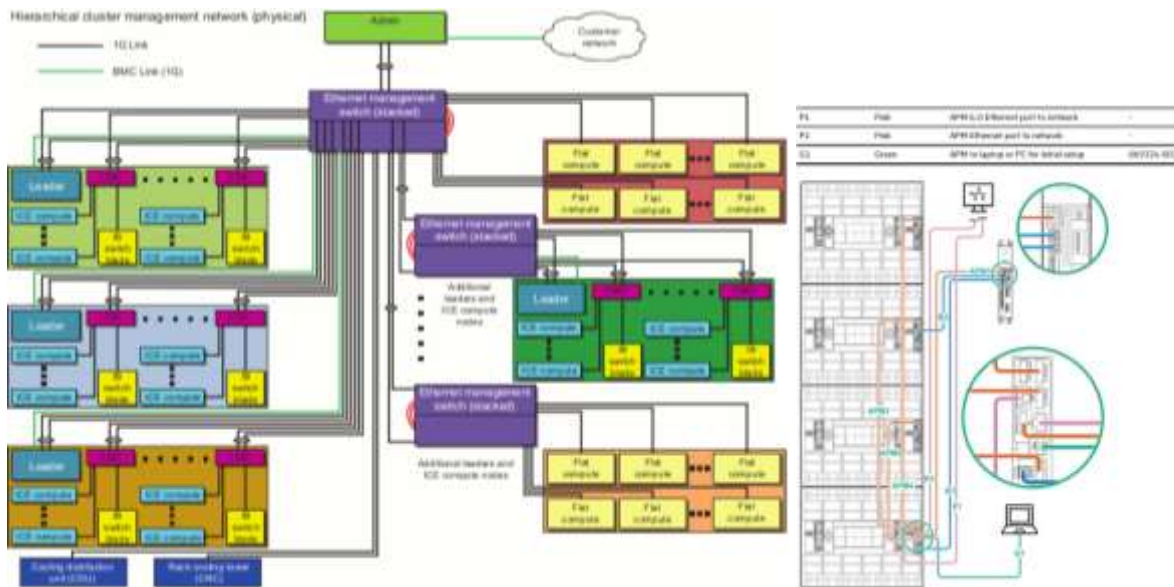
Enable Ethernet

Command

`ENABLE ETHERNET`

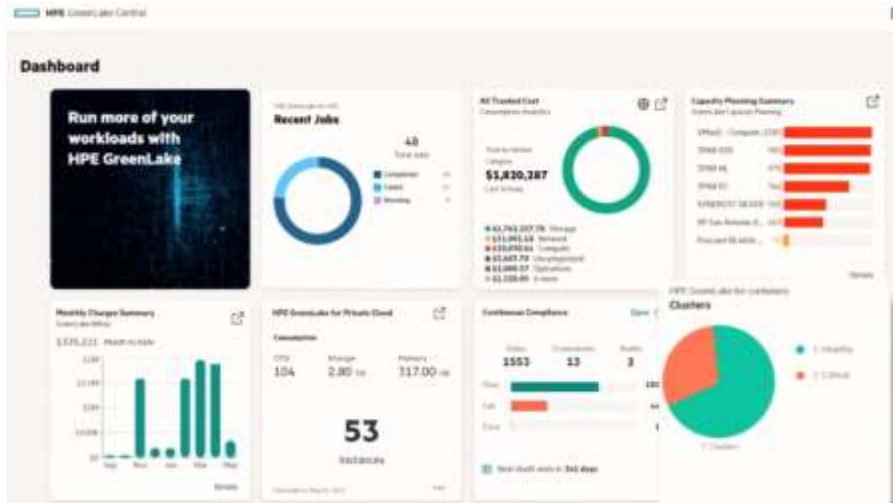
Description

Enables network access. Requires a password to be set and enabled SSH and Telnet access must be enabled or disabled separately. Also enables remote syslog support.



32. The Accused '841 Products further include a first cluster in a first configuration and a second cluster in a second configuration, each comprising a set of computing resources, including at least one hardware processor, wherein each of the first and second clusters is communicatively linked to the private communications network. For example, HPE GreenLake for HPC includes several cluster configurations, such as a “large memory” configuration and a “large CPU” configuration. Each configuration includes software and hardware specific to a particular use, such as, *e.g.*, HPC management and clustering software and HPE Apollo, HPE Cray, and/or HPE ProLiant computing/storage resources. Each compute resource includes a processor,

such as AMD EPYC 7003 Series processors, and is networked via a private communication network as noted above and illustrated below:



- Single or dual processor systems with AMD EPYC™ 7003 Series processors, including the power, frequency, or core count processors to match your workload requirements.

The screenshot shows the Clusters page in HPE GreenLake Central. It contains a table with the following columns: Name, Health, Size, App Health, Role, Updated Time, Status, CPU, and Memory. The table lists several clusters, all with a 'Healthy' status and 'Ready' role.

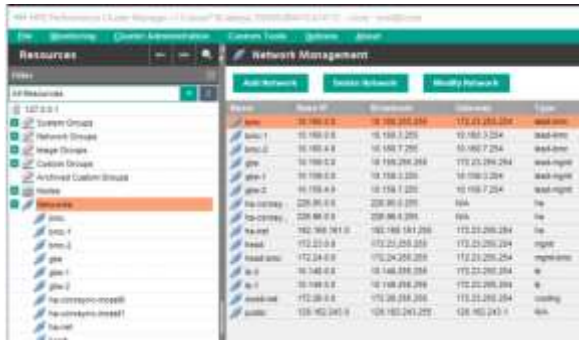
Name	Health	Size	App Health	Role	Updated Time	Status	CPU	Memory
hpegreenlakecluster	Healthy	Auto	Ready	Ready	2	Ready	100%	100%
hpegreenlakecluster	Healthy	Auto	Ready	Ready	2	Ready	100%	100%
hpegreenlakecluster	Healthy	Auto	Ready	Ready	2	Ready	100%	100%
hpegreenlakecluster	Healthy	Auto	Ready	Ready	2	Ready	100%	100%
hpegreenlakecluster	Healthy	Auto	Ready	Ready	2	Ready	100%	100%
hpegreenlakecluster	Healthy	Auto	Ready	Ready	2	Ready	100%	100%
hpegreenlakecluster	Healthy	Auto	Ready	Ready	2	Ready	100%	100%

The HPE Apollo and HPE Cray Systems supports a full complement of HPC options

Optimize HPC clusters for many different applications including product design and testing simulation, financial risk modeling and Monte Carlo simulation, and scientific research modeling.

HPC options include top bin CPUs, fast memory, integrated accelerators (GPUs or coprocessors), and fast cluster fabrics and I/O interconnections, making it easy for you to achieve the right performance, and price/performance for your HPC workloads.





Self-service for the provisioning of compute, storage, and networks is done through the HPE GreenLake Central console, and the entire HPC stack clusters, operating software and so on, is managed by HPE experts from one of a dozen centers around the world, customers operate the clusters with self-service capabilities in HPE GreenLake Central to manage queues, jobs, and output. HPE GreenLake for HPC gets HPC centers out of the business of maintaining the hardware and software of those clusters, and while HPE is not offering application management services, it does have a widget with self-service capabilities that will snap into the GreenLake Central console and it will entertain managing the HPC applications themselves under a separate contract if customers really want this.

Security features

The Linux distributions include security features. The cluster manager provides additional security features, and among those features are the following:

- Secure provisioning over the management network

Supported cluster blueprints

The following cluster blueprints are supported for use in HPE GreenLake for containers. All applications annotated with metadata will get a unique endpoint for customers to use to launch their applications.

large-memory

Blueprint to use when creating a high-availability, high performance cluster for applications such as the following:

- Java applications
- Database applications
- Multiprocess applications

large-cpu

Blueprint to use when creating a high-availability cluster that has three large-cpu worker nodes. This kind of cluster is useful for trying out CPU-intensive applications. This kind of cluster is recommended for applications such as the following:

- Jenkins
- Continuous integration and continuous development (CI/CD) applications
- Multiprocess applications
- Monitoring applications

Apollo 6000 System: rack-scale solutions with better density, performance, power efficiency, and cost of ownership


To address the growing demand for HPC and the relentless pursuit of efficiency, Hewlett Packard Enterprise has taken the lead on a new approach: thinking beyond just the server and designing a rack-level solution that gives you the right compute at the right economics so you can get the most out of your infrastructure—and your budget. The HPE Apollo 6000 System delivers industry-leading performance in less space with the flexibility to tailor the system to precisely meet workload requirements.

Apollo 6500 System: high density GPU compute

The HPE Apollo 6500 system solve problems faster with up to 15 Tflops of single-precision performance per 2U node. The new HPE Apollo 6500 increases your return on IT investment by accelerating the performance of your data center workloads with up to eight GPUs or coprocessors. With high-powered analysis and prediction, you will solve your most demanding problems in the shortest time.

33. In addition, the Accused '841 Products include a monitoring system monitoring operations of the first and second clusters, identifying operational and connectivity problems. For example, the HPE Accused Products include cluster level monitoring, e.g., cluster status, usage metrics, fabric health, and overall cluster health, which are presented via consoles and dashboards

such as, e.g., the HPE Performance Cluster Manager for HPC and GreenLake Central, as illustrated below:



Viewing details about HPE GreenLake for containers


The HPE GreenLake for containers card in the HPE GreenLake Central Dashboard shows information about your clusters.

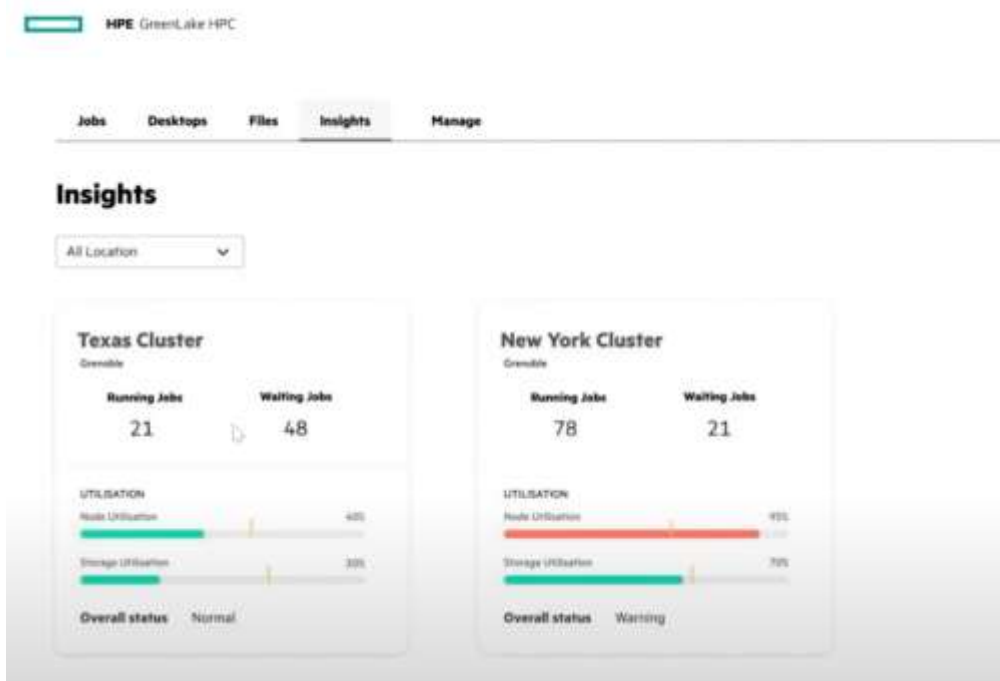
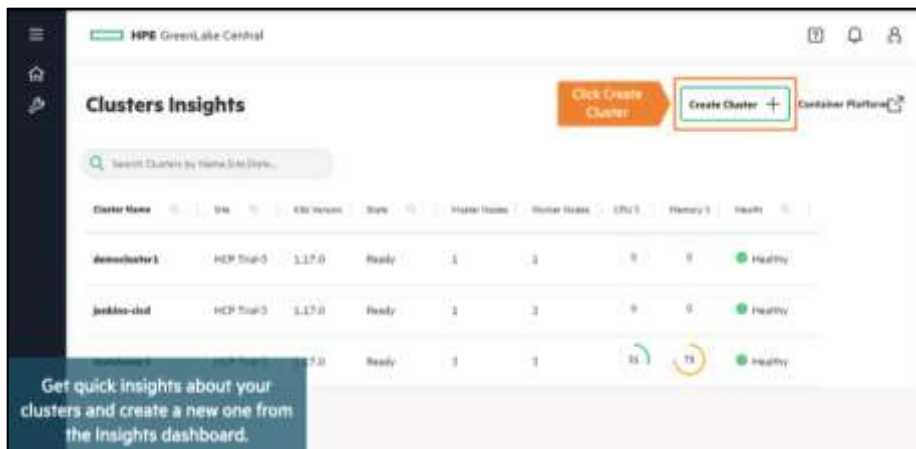
Prerequisites

- An active service subscription or trial subscription to HPE GreenLake for containers
- A user account in a tenant with access to the HPE GreenLake for containers service in HPE GreenLake Central

Procedure

1. Log in to HPE GreenLake Central at client.greenlake.hpe.com.
If you are already logged in, click **Dashboard** in the banner at the top of the screen.
2. Locate the HPE GreenLake for containers card.
3. Note the number and status of clusters created in the HPE Emserv Container Platform.
If no clusters have been created, you will see "Create Cluster" on the card.





34. In addition, the Accused '841 Products include a monitoring system issuing an alert in response to the identified problems indicating a corresponding one of the first and second clusters associated with the identified problems. For example, the HPE Accused Products monitor cluster-level metrics and display alerts with respect to such metrics in a way that identifies which alerts apply to which clusters, as illustrated below:




Cluster Health Dashboard - Cluster Status Report

Fabric Health Report

Cluster Health Report

Detailed Alerts

Hosts	Severity	Group	Resource	Resource	Resource	Event	Text
ok	Warning	cooling	ch01	cooling	cooling	water-making-level	Check Water Making Level Alarm
ok	Warning	cooling	ch02	cooling	cooling	facility-room-alert	FIR Facility Status Alarm
ok	Warning	cooling	ch04	cooling	cooling	water-making-level	Check Water Making Level Alarm
ok	Warning	cooling	ch04	cooling	cooling	water-making-empty	Water Making Empty Alarm
ok	Warning	cooling	ch01	cooling	cooling	water-making-level	Check Water Making Level Alarm
ok	Warning	cooling	ch02	cooling	cooling	water-making-empty	Water Making Empty Alarm
ok	Warning	cooling	ch02	cooling	cooling	water-making-level	Check Water Making Level Alarm
ok	Warning	cooling	ch02	cooling	cooling	gas-flue-alert	FSGIT Supply Pre-Flue Alarm
ok	Warning	cooling	ch02	cooling	cooling	airflow-alert	FRI IT Return Air Sensor Alarm
ok	Warning	cooling	ch02	cooling	cooling	server-alert	TIT Supply Air Sensor Alarm
nom	Warning	fabric	CR02T03A1	ibmc	net-port15	state-change	Switch (SWS-1224G0000170AR) Port 15 State is NOT active
down	Ok	fabric	CR02T03A1	ibmc	net-port15	link-speed-change	Switch Port 15 Link Speed is changed from 10000000 to 1000 Speed
down	Ok	fabric	CR02T03A1	ibmc	net-port15	link-width-change	Switch Port 15 Link Width changed from 10000000 to 10 Width
down	Ok	fabric	CR02T03A1	ibmc	net-port15	link-speed-change	Switch Port 3 Link Speed is changed from 10000000 to 1000 Speed
down	Ok	fabric	CR02T03A1	ibmc	net-port15	link-width-change	Switch Port 3 Link Width changed from 10000000 to 10 Width
down	Ok	fabric	CR02T03A1	ibmc	net-port15	state-change	Switch Port 3 State is LINKUP



Monitoring and cluster security

Monitoring consists of two conceptual parts:

- The back-end gathering and archiving of metrics and reactions to alerts
- The GUI metric and alert display

There are two methods of gathering metrics. The first method launches monitoring agents onto the compute nodes to gather in-band, operating system level metrics and aggregate them back to the admin node for storage and display. The second method supports invoking a program on the admin node. This program gathers metrics out-of-band, or outside of the running OS on the compute nodes, and feeds those metrics to the cluster manager for storage and display. An example of the out-of-band approach is the pre-configured support for gathering hardware metrics, such as power and temperature, from the iLO of each compute node.

The pie graphs in the global cluster view represent the cluster monitoring sensor value. To select the sensors being monitored, right-click an item in the right pane. A metric window displays. Select a metric and click OK.

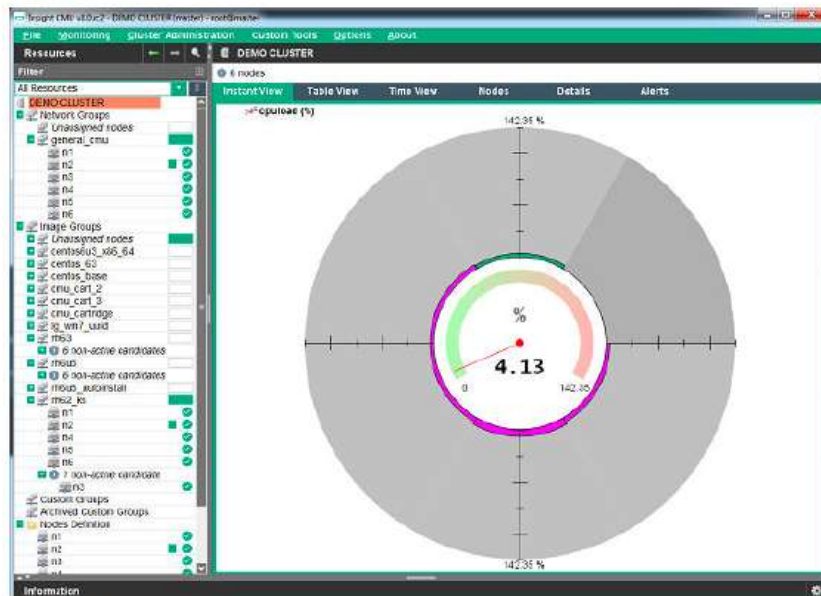


Figure 5: Monitoring window

35. In addition, the Accused '841 Products include clusters wherein the first configuration differs from the second configuration and wherein the first configuration provides a first computing environment for performing a first client task and the second configuration provides a second computing environment for performing a second client task. For example, the HPE Accused Products include cluster blueprints which vary in configuration depending on usage requirements, such as “large-memory” configurations or “large-cpu” configurations, each respectively used for differing tasks. Further, the Accused Products can be customized for specific applications, such as Big Data, ML/AI, *etc.*, as illustrated below:

Supported cluster blueprints

The following cluster blueprints are supported for use in HPE GreenLake for containers.

All applications annotated with metadata will get a unique endpoint for customers to use to launch their applications.

large-memory

Blueprint to use when creating a high-availability, high performance cluster for applications such as the following:

- Java applications
- Database applications
- Multiprocess applications

Object	Value
Kubernetes version	1.17.4 ¹
HA (high availability)	Yes
Master nodes	3 nodes: <code>hpe-k8s-standard-master</code> blueprint
Worker Nodes	3 nodes: <code>hpe-k8s-large-mem-worker</code> blueprint

large-cpu

Blueprint to use when creating a high-availability cluster that has three large-cpu worker nodes. This kind of cluster is useful for trying out CPU-intensive applications. This kind of cluster is recommended for applications such as the following:

- Jenkins
- Continuous integration and continuous development (CI/CD) applications
- Multiprocess applications
- Monitoring applications

Object	Value
Kubernetes version	1.17.4 ¹
HA (high availability)	Yes
Master nodes	3 nodes: <code>hpe-k8s-standard-master</code> blueprint
Worker Nodes	3 nodes: <code>hpe-k8s-large-cpu-worker</code> blueprint

The Flemish Tier 1 supercomputer opened at Ghent University in October 2012, and is operated by Ghent University and shared by all Flemish universities and by industry and public research institutions. The Tier 1 supercomputer facility consists of a 528-node cluster of HPE ProLiant SL230s Gen8 Servers, and it also includes HPE ProLiant DL380e Gen8 Servers.

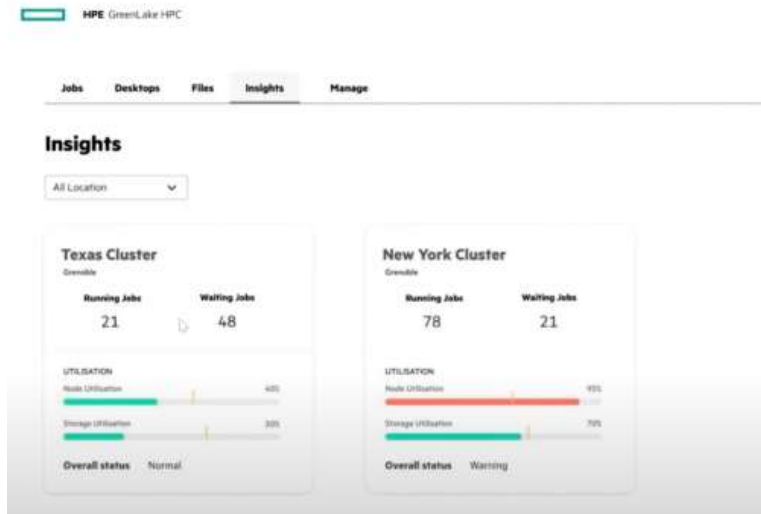
The Apollo 6000 System delivers industry-leading performance in less space with the flexibility to tailor the system to precisely meet workload requirements. This rack-level solution allows Ghent University to address the growing demand for high-performance computing while controlling costs and optimizing operational efficiency. It allows Ghent University to optimize performance at the rack level, and the HPE Apollo family is designed to deliver efficient rack-scale compute, storage, networking, power, and cooling solutions for high-performance computing workloads.

With over 41,000 students and 9,000 employees, Ghent University is one of the largest universities in Belgium, and it sought additional high-performance computing power to support its own researchers and therefore extended its Tier 2 supercomputing cluster to strengthen the university's research capabilities. After a careful evaluation, Ghent University selected the HPE Apollo 6000 System and is now able to conduct even more sophisticated computational research than previously possible.

The Apollo 6000 System was deployed with ProLiant XL230a Gen9 Servers and runs Red Hat Enterprise Linux 7. Twenty Apollo 6000 System chassis were deployed with 10 nodes per chassis in a scaled-out chassis configuration. In addition, 16 ProLiant DL380e Gen9 Servers were each fitted with half a TB of memory and a tiered HD/SSD storage solution in order to boost Big Data research. The installation was smooth, and the new computing cluster went online in April 2015. The academic community immediately began reaping the rewards of the expanded high-performance computing environment.

Harnessing data growth with purpose-built HPC and AI solutions

The new service includes a combination of HPE Apollo systems and HPE ProLiant servers, which ingest and process high volumes of data, and support deep learning and artificial intelligence capabilities. As part of the HPE GreenLake service, HPE will build and manage the complete solution that will be hosted at a QTS data center, a hosting facility that delivers secure, compliant data center infrastructure and robust connectivity to support scaling of operations.





Apollo 6000 System: rack-scale solutions with better density, performance, power efficiency, and cost of ownership

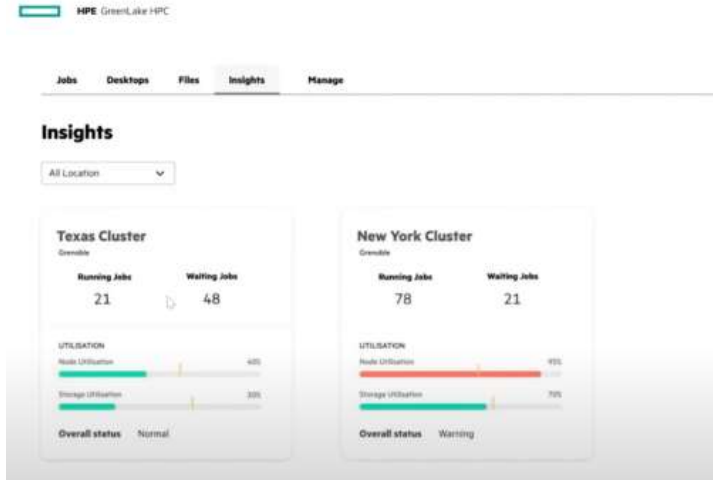
To address the growing demand for HPC and the relentless pursuit of efficiency, Hewlett Packard Enterprise has taken the lead on a new approach: thinking beyond just the server and designing a rack-level solution that gives you the right compute at the right economics so you can get the most out of your infrastructure—and your budget. The HPE Apollo 6000 System delivers industry-leading performance in less space with the flexibility to tailor the system to precisely meet workload requirements.

Apollo 6500 System: high density GPU compute

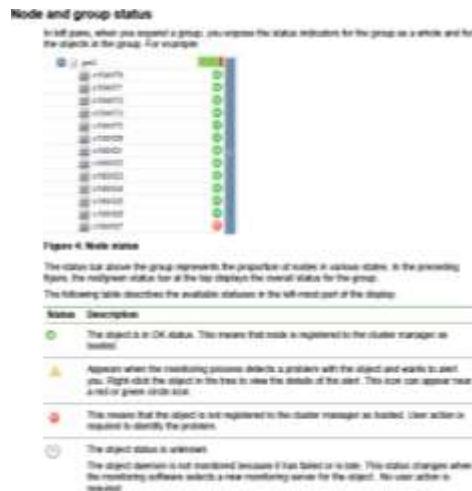
The HPE Apollo 6500 system solve problems faster with up to 15 Tflops of single-precision performance per 2U node. The new HPE Apollo 6500 increases your return on IT investment by accelerating the performance of your data center workloads with up to eight GPUs or coprocessors. With high-powered analysis and prediction, you will solve your most demanding problems in the shortest time.

36. Furthermore, the Accused '841 Products include a monitoring system comprising a main monitor that operates to monitor the first and second clusters to identify the operation and connectivity problems. For example, the Accused Products monitor cluster status, cluster-level alerts, cluster fabric, and connectivity and consumption information for each cluster and identify operational or network problems, such as low resource capacity or network fabric errors, as seen below:





37. Additionally, the Accused '841 Products include a monitoring system comprising monitors for each node of the first and second clusters operating to check for hardware and software problems within a particular node and to report the hardware and software problems to the main monitor. For example, the Accused Products provide status indicators for clusters and groups as a whole and for objects within the group, such as each individual node, and a node view functionality including displaying of alerts regarding performance issues, as seen below:



To display details for a specific node, select the node in the left-panel tree. This displays the Node view in the central frame.

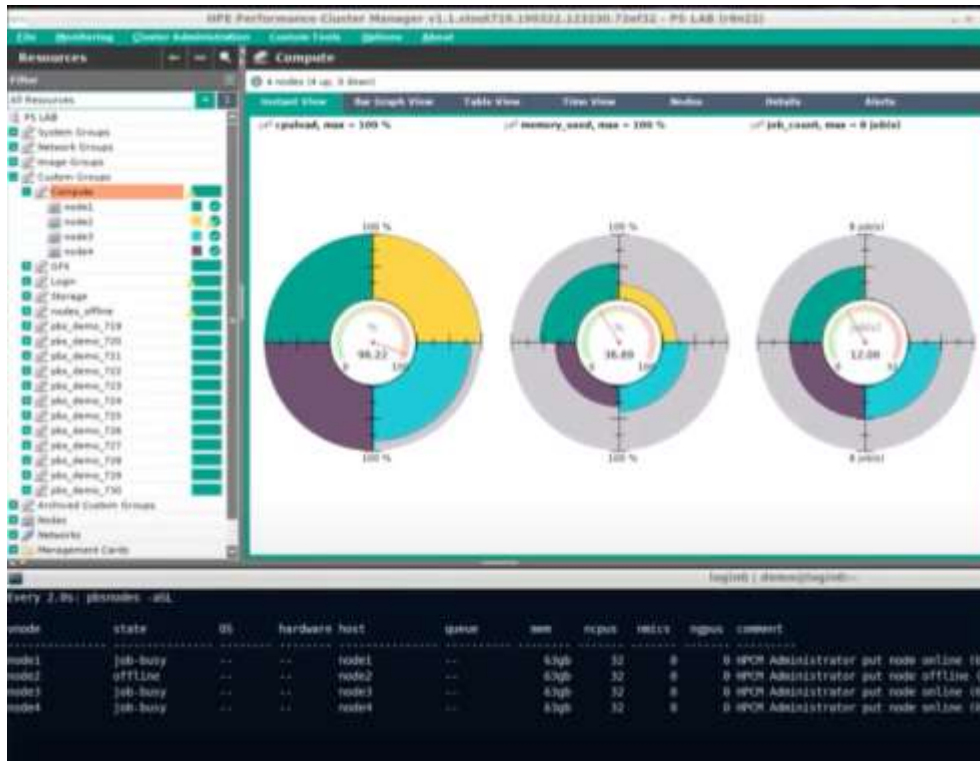
The following tabs are available in Node view:

- **Monitoring** - Displays monitoring metric values for the node.
- **Details** - Displays static data for the node. Some values are filled during the initial node discovery (scan node). For others, right-click the node in the left-panel tree and select **Update > Update Node Details**.
- **Alerts** - Contains the alerts currently raised for the specified node.

The terminal window displays system information for a node named 'node1'. The output is organized into sections: 'General Information', 'Network Information', 'System Information', and 'Hardware Information'. Key details include the node name 'node1', IP address '100.138.1.100', and various system parameters like 'Platform name' and 'Architecture'.

```
ALERTS
#
#cpu_freq_alert "CPU frequency is not nominal" 1 24 100
< & ah -c "bc"cat /sys/devices/system/cpu/cpu0/cpufreq/
scaling_cur_freq;echo 100 \* \b / \a |bc"
login_alert "Someone is connected" 3 24 0 >
login(s) w -h | wc -l
root_fa_used "The / filesystem is above 90% full" 4 24
90 > & df / | awk '{ if ($6=="") print $5}' | cut -f 1 -d
\ -
#reboot_alert "Node rebooted" 4 24 5 < rebooted awk
'(printf "%i,%i\n",01/60)' /proc/uptime
# The line below allows to report MCE errors; be careful for possible false
positives
#mce_alert "The kernel has logged MCE errors; please check /var/log/
mcelog" 5 60 i > lines wc -l /var/log/mcelog |cut -f 1 -d ' '
```

The screenshot shows the 'Insights' section of the HPE GreenLake HPC dashboard. It features a navigation bar with 'Jobs', 'Desktops', 'Files', 'Insights', and 'Manage'. The 'Insights' section includes a dropdown for 'All Location' and two cluster overview cards: 'Texas Cluster' and 'New York Cluster'. Each card displays 'Running Jobs' and 'Waiting Jobs' counts, along with 'Node Utilization' and 'Storage Utilization' progress bars. The Texas Cluster shows 21 running and 48 waiting jobs, with an overall status of 'Normal'. The New York Cluster shows 78 running and 21 waiting jobs, with an overall status of 'Warning'.



38. Additionally, HPE has been, and currently is, an active inducer of infringement of the '841 patent under 35 U.S.C. § 271(b) and a contributory infringer of the '841 patent under 35 U.S.C. § 271(c).

39. HPE has actively induced, and continues to actively induce, infringement of the '841 patent by intending that others use, offer for sale, or sell in the United States, products and/or services covered by one or more claims of the '841 patent, including but not limited to, the HPE GreenLake for High Performance Computing (HPC), HPE GreenLake for Containers, GreenLake for Private Cloud, HPC Performance Cluster Manager, and related HPE Apollo and ProLiant Systems, as well as any HPE product and/or service, alone or in combination, that operates in materially the same manner. HPE provides these products and/or services to others, such as customers, resellers and end-user customers, who, in turn, in accordance with HPE's design, intent and directions, use, provision for use, offer for sale, or sell in the United States the foregoing products and/or services that directly infringe one or more claims of the '841 patent as described

above. HPE's inducement includes the directions and instructions found at one or more of the following, the content of which is specifically illustrated above:

- https://www.hpe.com/psnow/doc/a00093985en_us
- <https://www.hpe.com/us/en/newsroom/press-release/2021/09/hewlett-packard-enterprise-wins-2b-hpe-greenlake-contract-with-the-national-security-agency.html>
- <https://support.hpe.com/hpesc/public/docDisplay?docId=c04558867>
- <https://www.hpe.com/psnow/doc/4AA5-8958ENW.pdf>
- https://support.hpe.com/hpesc/public/docDisplay?docId=a00069799en_us
- https://www.hpe.com/us/en/discover.html?media-id=%2Fus%2Fen%2Fresources%2Fdiscover%2Fflas-vegas-2021%2FSL4368%2F_jcr_content.details.json
- <https://youtu.be/SZPUb9-cQ>
- <https://www.youtube.com/watch?v=Z8PdC8fXJc8>
- <https://www.youtube.com/watch?v=BAHyXHwbQEE>
- <https://www.hpe.com/us/en/greenlake/containers.html#demo>
- https://support.hpe.com/hpesc/public/docDisplay?docId=a00048260en_us&docLocale=en_US
- https://www.hpe.com/psnow/doc/a00092451en_us
- https://support.hpe.com/hpesc/public/docDisplay?docId=a00048260en_us&docLocale=en_US
- <https://www.youtube.com/watch?v=KLhi76yODPo>
- HPE Ezmeral Contain Platform 5.1 Documentation (including documentation for later versions as well)

40. HPE has contributed to, and continues to contribute to, the infringement of the '841 patent by others by knowingly providing one or more components that, when installed, configured, and used result in systems that, as intended by HPE described above, directly infringe one or more claims of the '841 patent.

41. HPE knew of the '841 patent, or should have known of the '841 patent, but was willfully blind to its existence. Upon information and belief, HPE has had actual knowledge of

the '841 patent since at least as early as the receipt of IV's December 14, 2021 notice letter, which attached a copy of the '841 patent, and service upon HPE of the Complaint in this case.

42. By the time of trial, HPE will or should have known and intended (since receiving such notice) that its continued actions would infringe and actively induce and contribute to the infringement of one or more claims of the '841 patent.

43. HPE has committed, and continues to commit, contributory infringement by, *inter alia*, knowingly selling products and/or services that when used cause the direct infringement of one or more claims of the '841 patent by a third party, and which have no substantial non-infringing uses, or include one or more separate and distinct components such as software especially made or adapted for use in infringement of the '841 patent that are not staple articles or commodities of commerce suitable for substantial non-infringing use.

44. As a result of HPE's acts of infringement, IV has suffered and will continue to suffer damages in an amount to be paid at trial.

PRAYER FOR RELIEF

IV requests that the Court enter judgment as follows:

- (A) that HPE has infringed one or more claims of the asserted patent, directly and/or indirectly, literally and/or under the doctrine of equivalents;
- (B) awarding damages sufficient to compensate IV for HPE's infringement under 35 U.S.C. § 284;
- (C) finding this case exceptional under 35 U.S.C. § 285 and awarding IV its reasonable attorneys' fees;
- (D) awarding IV its costs and expenses incurred in this action;
- (E) awarding IV prejudgment and post-judgment interest; and

(F) granting IV such further relief as the Court deems just and appropriate.

DEMAND FOR JURY TRIAL

IV demands trial by jury of all claims so triable under Federal Rule of Civil Procedure 38.

Date: December 15, 2021

Respectfully submitted,

/s/ Derek Gilliland

DEREK GILLILAND

State Bar No. 24007239

SOREY & GILLILAND, LLP

133 E. Tyler Street

Longview, Texas 75601

903.230.5600 (telephone)

derek@soreylaw.com

OF COUNSEL:

Paul J. Hayes

phayes@princelobel.com

Matthew Vella

mvella@princelobel.com

Robert R. Gilman

rgilman@princelobel.com

Jonathan DeBlois

jdeblois@princelobel.com

Brian Seeve

bseeve@princelobel.com

Thomas Fulford

tfulford@princelobel.com

PRINCE LOBEL TYE LLP

One International Place, Suite 3700

Boston, MA 02110

Tel: (617) 456-8000

Fax: (617) 456-8100

ATTORNEYS FOR PLAINTIFF

INTELLECTUAL VENTURES II LLC