


SOLUTION BRIEF

Enterprise Data Center Infrastructure
2nd Generation Intel® Xeon® Scalable Processors
April 2019



Intel® Reference Solution for IBM Cloud Private for Data*

Accelerate your path to an enterprise-grade analytics platform with an optimized, verified solution for VMware vSAN*.

Digital transformation isn't a buzzword; it's a fact of life for organizations today, and artificial intelligence (AI) is a significant consideration. As an example, according to IDC, by 2022, over 50 percent of IT security alerts will be handled by AI.¹ Effective, AI-based digital transformation needs to be at the center of corporate strategy for enterprise customers to garner business insights that enable them to stay ahead of the competition.

This transformation brings new priorities. For example, 81 percent of businesses report AI to be an important area for investment.² However, AI brings with it new data challenges, particularly for training AI models. Beyond the amount and diversity of data, organizations struggle with the diversity of data *platforms*. Piecemeal tools produce piecemeal results; enterprises need an integrated data science platform to meet the data and analytics challenges posed by digital transformation.

IBM Cloud Private for Data* provides an integrated, extensible analytics platform with all the tools and services necessary to navigate multiple clouds, effectively harness data, and extract value using powerful analytics and AI tools. It is built on IBM Cloud Private*, an enterprise-grade private cloud that is reliable, efficient, scalable, and portable.

The combination of IBM Cloud Private for Data, VMware vSphere*, and VMware vSAN* on Intel® hardware can provide an economical, enterprise-grade way forward for digital transformation. The Intel® Reference Solution for IBM Cloud Private for Data combines 2nd Generation Intel® Xeon® Scalable processors, Intel® 3D NAND Solid State Drives (SSDs), Intel® Optane™ DC SSDs, and the Intel® Ethernet 700 Series with the multi-cloud analytics capabilities of IBM Cloud Private for Data and the hyper-converged storage offered by VMware vSAN to enable businesses to quickly deploy an enterprise-ready, extensible analytics platform with confidence.

Authors

Rick Cossen

Solution Director
Intel Datacenter AI Solutions
Development Group

Raghu K. Moorthy

Software Alliances Director,
Intel Data Center Sales Group

Lokendra Uppuluri

Solutions Architect,
Intel Datacenter AI Solutions
Development Group

IBM Cloud Private for Data

IBM Cloud Private for Data is a market-leading, integrated analytics platform that is built on IBM Cloud Private and harnesses the power of Intel hardware and technologies. It provides an integrated data-management and analytics platform that enables you to:



Manage All
Your Data



Build Your Ladder
to AI



Modernize Your
Data and Analytics
Workloads



Improve Your
Compliance
Readiness

Figure 1. Key capabilities of IBM Cloud Private for Data*

- **Manage all your data**—provide governed, more-secure access to your data regardless of where it lives
- **Build your ladder to AI**—deploy the IT infrastructure you need to enable self-service analytics and AI
- **Modernize your data and analytics workloads**—shift from monolithic applications to microservices to increase agility and efficiency
- **Improve your compliance readiness**—accelerate your transformation to meet new regulatory requirements like the EU's General Data Protection Regulation (GDPR)

IBM Cloud Private for Data Benefits for Analytics

The capabilities of IBM Cloud Private for Data provide a number of key benefits for your analytics workloads.

Multi-cloud Support

By building in multi-cloud support at its foundation, IBM Cloud Private for Data enables customers to build once, run anywhere. IBM Cloud Private for Data externalizes AI models and applications as Kubernetes* services and enables organizations to run analytics and AI workloads in their own data centers or on the cloud infrastructures of their choice. This flexibility is central to organizational agility; it enables organizations to provide continuous delivery of analytics and AI services and to quickly adapt to changing industry partnerships and opportunities.

Data Virtualization

Many companies are not unlocking the value they would like from their data; most struggle to manage an average of 33 unique data sources, often trapped in data silos that are hard to find and access.³ Data virtualization in IBM Cloud Private for Data enables users to query data across many systems without having to copy and replicate data, which helps reduce costs. It also simplifies analytics and makes companies' analytics more up-to-date and accurate because users query the latest data at its source.

Operationalizing Analytics and AI

Analogous to DevOps, AnalyticOps capabilities in IBM Cloud Private for Data provide a collaborative environment that bridges the gap between analytics and their operational deployment at all phases of the AI lifecycle: data collection and preparation, model training and testing, and deployment and monitoring.

IBM Cloud Private for Data provides AI collaboration across all roles in an organization through IBM Watson OpenScale*, which enables organizations to monitor how deployed AI models affect business outcomes, regardless of where models are developed or deployed.

Data Security

Data security is built into every part of IBM Cloud Private for Data. The solution supports data encryption through the storage layer and makes use of the underlying security, logging, and monitoring capabilities of Intel hardware and IBM Cloud Private. Authentication and authorization is built into IBM Cloud Private for Data. It integrates with the industry-leading Lightweight Directory Access Protocol (LDAP), such as is used by Microsoft Active Directory*, and it also supports Security Assertion Markup Language (SAML) for token-based single sign-on (SSO). IBM Cloud Private for Data also supports data privacy and governance by design and includes data masking, auto-discovery of data, and data curation.

VMware vSAN ReadyNodes*

The Intel Reference Solution for IBM Cloud Private for Data uses VMware vSAN ReadyNodes* to deliver hyper-converged infrastructure and serve as the foundation for a transformed, software-defined data center. The VMware vSAN ReadyNode certification program provides assurance to data center buyers that their VMware vSAN provider of choice has undergone VMware's rigorous certification process. Verified Intel solutions—such as this reference solution—are certified for vSAN ReadyNodes and are tightly specified by Intel and VMware to deliver out-of-the-box high performance. With the Intel Reference Solution for IBM Cloud Private for Data, IT organizations can rest assured that their solutions are already verified for balanced and optimized performance, from the hardware up through the firmware stack to the VMware vSAN software. IT teams can get right to work providing VMware vSAN services to IBM Cloud Private for Data, rather than wading through multiple component options or conducting extensive, system-level testing.

Intel Reference Solution for IBM Cloud Private for Data

The Intel Reference Solution for IBM Cloud Private for Data helps optimize price and performance, and it can reduce infrastructure evaluation time compared to building an IBM Cloud Private for Data solution from scratch. Specifically, the reference solution combines Intel data center processors, storage, and network hardware to empower enterprises to quickly harness a reliable, comprehensive solution that allows organizations to:

- **Protect** infrastructure investments with scalable, general-purpose storage and compute
- **Generate excellent total cost of ownership (TCO)** with industry-standard hardware that IT organizations are used to managing
- **Accelerate time to market** by using a solution with a rich development toolset that is optimized for crucial software libraries for analytics, AI, and data science

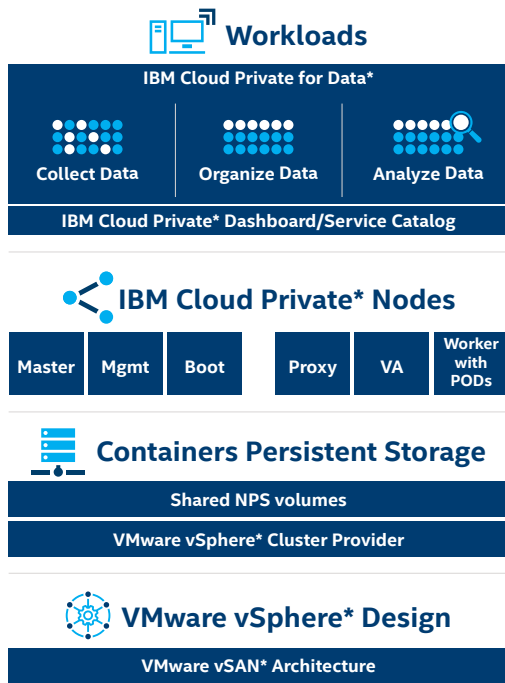


Figure 2. Software design for IBM Cloud Private for Data* on VMware vSAN*

Solution Components

The Intel Reference Solution for IBM Cloud Private for Data combines 2nd Generation Intel Xeon Scalable processors, Intel 3D NAND SSDs, Intel Optane DC SSDs, and the Intel Ethernet 700 Series, all of which enable businesses to quickly deploy IBM Cloud Private for Data on VMware vSAN on a performance-optimized infrastructure.

2nd Generation Intel® Xeon® Scalable Processors

2nd Generation Intel Xeon Scalable processors provide the Intel Reference Solution for IBM Cloud Private for Data with an excellent performance-to-cost ratio. Intel Xeon Platinum processors, part of the Intel Xeon Scalable processor family, are also available for the solution and they can provide extra performance when required.

Intel® SSD Data Center Family

Storage latency can be a bottleneck for container and virtual machine (VM) performance. For this reason, the Intel Reference Solution for IBM Cloud Private for Data uses Intel Optane SSD DC P4800X and Intel SSD DC P4510 drives. Based on Intel Optane technology and Intel 3D NAND technology, these enterprise data center SSDs can provide 13x better performance than hard-disk drives (HDDs).⁴

Intel® Ethernet Connections and Intel® Ethernet Adapters

The Intel Ethernet 700 Series accelerates the performance of the Intel Reference Solution for IBM Cloud Private for Data. The Intel Ethernet 700 Series delivers validated performance ready to meet high-quality thresholds for data resiliency and service reliability with broad interoperability.⁵ All Intel Ethernet products are backed by worldwide pre- and post-sales support and offer a limited lifetime warranty.

2nd Generation Intel® Xeon® Scalable Processors

2nd Generation Intel Xeon Scalable processors:

- Offer high scalability that is cost-efficient and flexible, from the multi-cloud to the intelligent edge
- Establish a seamless performance foundation to help accelerate data's transformative impact
- Support breakthrough Intel® Optane™ DC persistent memory technology
- Accelerate artificial-intelligence (AI) performance and help deliver AI readiness across the data center
- Provide hardware-enhanced platform protection and threat monitoring

The Intel® Reference Solution for IBM Cloud Private for Data* features 2nd Generation Intel Xeon Silver processors, 2nd Generation Intel Xeon Gold processors, and 2nd Generation Intel Xeon Platinum processors.



Dev/Test, Base, Plus, and Custom Configurations

The Intel Reference Solution for IBM Cloud Private for Data is available in four configurations: “Dev/Test,” “Base,” “Plus,” and “Custom,” as shown in Appendix B. The Dev/Test configuration specifies the minimum required performance capability for the solution. The Base configuration offers an optimized balance of price and performance for a typical deployment. The Plus configuration provides one example of how system builders, system integrators, and solution and service providers can further optimize the solution to achieve higher performance and capabilities. The Custom configuration provides an example of optimizing the solution for large clusters.

Transform Your Infrastructure Faster with the Intel Reference Solution for IBM Cloud Private for Data

IBM and Intel modernize the data center for digital transformation. The combination of IBM Cloud Private for Data and VMware vSAN provides an open, agile, and scalable platform for building modern multi-cloud workloads. Coupled with Intel hardware and technologies, this provides enterprise-grade security and performance for multi-cloud workloads—on the broadest software platform—to give you choice with consistency in building multi-cloud environments. When organizations choose the Intel Reference Solution for IBM Cloud Private for Data, they get pre-tuned and tested configurations that are workload-optimized and that let organizations deploy data center infrastructure quickly and efficiently with less tuning.

Learn More

1. IBM Cloud Private for Data: ibm.com/products/cloud-private-for-data
2. Intel and IBM: ibm.com/cloud/private/partners
3. Intel Xeon Scalable processors: intel.com/xeonscalable
4. Intel SSD Data Center Family: intel.com/content/www/us/en/products/memory-storage/solid-state-drives/data-center-ssds.html
5. Intel Ethernet 700 Series: intel.com/ethernet

Appendix A: Quantity of Physical Nodes for the Intel Reference Solution for IBM Cloud Private for Data

The IBM Cloud Private for Data cluster at the heart of the Intel Reference Solution for IBM Cloud Private for Data has four main classes of nodes: boot, master, worker, and proxy. In addition, it has two optional classes of nodes—management and vulnerability advisor. The chart below lists the number of physical nodes by hardware configuration of the solution that are required to support a given range of VM-based workloads. For additional details on VM sizing, sample workload, and cluster verification, see the Intel Reference Solution for IBM Cloud Private for Data reference architecture.

CONFIGURATION	SMALL IBM CLOUD PRIVATE FOR DATA*	MEDIUM IBM CLOUD PRIVATE FOR DATA	LARGE IBM CLOUD PRIVATE FOR DATA
APPROXIMATE NUMBER OF VIRTUAL MACHINES	6 WORKER-NODE VMS + 1 DATA-VIRTUALIZATION ADD-ON WORKER-NODE VM	12 WORKER-NODE VMS + 1 DATA-VIRTUALIZATION ADD-ON WORKER-NODE VM + 1 IBM DB2* ADD-ON WORKER-NODE VM	18 WORKER-NODE VMS + 1 DATA-VIRTUALIZATION ADD-ON WORKER-NODE VM + 2 DB2 ADD-ON WORKER-NODE VMS
USERS SUPPORTED	40	80	120
DEV/TEST SKU	7	N/A	N/A
BASE SKU	5	8	10
PLUS SKU	4	6	8
CUSTOM SKU	N/A	5	7

The configurations described above include the capacity needed for boot, master, proxy, management, and vulnerability-advisor node VMs. The number of users specified assumes that 25 percent of the users are generic analytics users, 25 percent are data scientists, 35 percent are data engineers, 14 percent are data stewards, and 1 percent are admins.

Appendix B: Dev/Test, Base, Plus, and Custom Configurations for the Intel Reference Solution for IBM Cloud Private for Data

Server vendor or data center solution providers providing this solution must meet or exceed the defined minimum configuration ingredients and reference minimum benchmark-performance thresholds listed below for this solution.

INGREDIENT	INTEL® REFERENCE SOLUTION FOR IBM CLOUD PRIVATE FOR DATA* DEV/TEST CONFIGURATION	INTEL REFERENCE SOLUTION FOR IBM CLOUD PRIVATE FOR DATA BASE CONFIGURATION	INTEL REFERENCE SOLUTION FOR IBM CLOUD PRIVATE FOR DATA PLUS CONFIGURATION	INTEL REFERENCE SOLUTION FOR IBM CLOUD PRIVATE FOR DATA CUSTOM CONFIGURATION
HARDWARE				
PROCESSOR	2 x Intel® Xeon® Silver 4210 processor (2.20 GHz, 10 cores, 20 threads), or a higher number Intel Xeon Scalable processor	2 x Intel Xeon Gold 5218 processor (2.30 GHz, 16 cores, 32 threads), or a higher number Intel Xeon Scalable processor	2 x Intel Xeon Gold 6248 processor (2.50 GHz, 20 cores, 40 threads), or a higher number Intel Xeon Scalable processor	2 x Intel Xeon Platinum 8268 processor (2.90 GHz, 24 cores, 48 threads), or a higher number Intel Xeon Scalable processor
MEMORY	192 GB or higher (12 x 16 GB DDR4-2,666 MHz)	384 GB or higher (12 x 32 GB DDR4-2,666 MHz)	768 GB** or higher (24 x 32 GB DDR4-2,666 MHz)	768 GB** or higher (24 x 32 GB DDR4-2,666 MHz)
BOOT DRIVE	1 x 480 GB Intel® SSD DC S3520 (M.2 Serial ATA [SATA])	2 x 480 GB Intel SSD DC S3520 (M.2 SATA)	2 x 480 GB Intel SSD DC S3520 (M.2 SATA)	2 x 480 GB Intel SSD DC S3520 (M.2 SATA)

Solution Brief | Intel® Reference Solution for IBM Cloud Private for Data*

CACHE TIER	1 x 375 GB Intel® Optane™ SSD DC P4800X (2.5-in. PCIe*)	2 x 375 GB Intel Optane SSD DC P4800X (2.5-in. PCIe)	2 x 375 GB Intel Optane SSD DC P4800X (2.5-in. PCIe)	2 x 375 GB Intel Optane SSD DC P4800X (2.5-in. PCIe)
CAPACITY TIER	3 x 2 TB Intel SSD DC P4510 (2.5-in. PCIe 3.1)	6 x 2 TB Intel SSD DC P4510 (2.5-in. PCIe 3.1)	8 x 2 TB Intel SSD DC P4510 (2.5-in. PCIe 3.1)	10 x 2 TB Intel SSD DC P4510 (2.5-in. PCIe 3.1)
DATA NETWORK	1 x Intel® Ethernet Converged Network Adapter X710-DA2	1 x Intel Ethernet Converged Network Adapter X710-DA2	1 x Intel Ethernet Converged Network Adapter X710-DA4 OR 2 x 25Gb Intel® Ethernet Controller XXV710	2 x 25Gb Intel Ethernet Controller XXV710
MANAGEMENT NETWORK PER NODE	Integrated 1 gigabit Ethernet (GbE)	Integrated 1 GbE	Integrated 1 GbE	Integrated 1 GbE
ADDITIONAL COMPONENTS		2 x Intel® 8-Port PCIe Gen3 x8 Switch AIC (AXXP3SWX08080)	2 x Intel 8-Port PCIe Gen3 x8 Switch AIC (AXXP3SWX08080)	2 x Intel 8-Port PCIe Gen3 x8 Switch AIC (AXXP3SWX08080)


SOFTWARE

IBM CLOUD PRIVATE*	Enterprise Edition 3.1.2
IBM CLOUD PRIVATE FOR DATA	Enterprise Edition 1.2.1.1 (CC11ML)
IBM DB2 WAREHOUSE*	3.3.0 (CCOEMEN)
DATA-VIRTUALIZATION ADD-ON	1.2.1.1 (CC11REN)
LINUX* OS	Red Hat* Enterprise Linux (RHEL*) 7.5
VMWARE VSPHERE*	6.7 U1
VMWARE VCENTER* APPLIANCE	6.7 U1
VMWARE VSAN*	6.7 U1

APPLIES TO ALL NODES

TRUSTED PLATFORM MODULE (TPM)	TPM 2.0 discrete or firmware TPM (Intel® Platform Trust Technology [Intel® PTT]) enabled
FIRMWARE AND SOFTWARE OPTIMIZATIONS	Intel® Trusted Execution Technology (Intel® TXT) enabled Intel® Hyper-Threading Technology (Intel® HT Technology) enabled Intel® Turbo Boost Technology enabled Intel® Speed Shift Technology, hardware P-states (HWP) native enabled Intel® Volume Management Device (Intel® VMD) disabled Power-management settings set to performance, workload input/output (I/O) intensive LLC prefetch enabled Uncore frequency scaling enabled PCIe and Single-Root I/O Virtualization (SR-IOV) support enabled Integrated I/O with Intel® Virtualization Technology for Directed I/O (Intel® VT for Directed I/O) enabled

**A cost-optimized configuration can use 512 GB RAM (2 x CPU, 16 x 32 GB 2,666 MHz DDR4). However, this is an unbalanced configuration and will decrease memory-access performance.



¹ IDC. "IDC FutureScape: Multiplied Innovation Takes Off, Powered by AI, Distributed Public Cloud, Microservices, Developer Population Explosion, Greater Specialization and Verticalization, and Scaling Trust." October 2018. idc.com/getdoc.jsp?containerId=prUS44417618.

² Fortune. "Fortune 500 CEOs on Trump, the Economy, and Artificial Intelligence." June 2017. <http://fortune.com/2017/06/08/fortune-500-companies-ceo-survey/>.

³ IBM. "Data virtualization." ibm.com/analytics/data-virtualization.

⁴ Evaluator Group. "Latest Intel Technologies Power New Performance Levels on VMware vSAN – 2018 Update." August 20, 2018. evaluatorgroup.com/document/lab-insight-latest-intel-technologies-power-new-performance-levels-vmware-vsan-2018-update/. 13x performance claim based on IOmark-VM* benchmark. Prior generation: 2 x Intel® Xeon® processor E5-2699 v4 (22 cores at 2.20 GHz with Intel® Hyper-Threading Technology [Intel® HT Technology] enabled) running VMware ESXi 6.0* with VMware vSAN 6.2*, 1 x Intel® SSD DC S3700 + 4 x 1 TB Seagate 10K HDD*; 88 IOmark-VM score. Current generation: 2 x Intel Xeon Gold 6154 processor (24 cores at 2.70 GHz with Intel HT Technology enabled) running VMware ESXi 6.7 with vSAN 6.7, 2 x Intel® Optane™ SSD DC P4800X + 4 x Intel SSD DC P4500, no deduplication; 1,152 IOmark-VM score.

⁵ The Intel® Ethernet 700 Series includes extensively tested network adapters, accessories (optics and cables), hardware, and software, in addition to broad operating system support. A full list of the product portfolio's solutions is available at intel.com/ethernet. Hardware and software is thoroughly validated across Intel® Xeon® Scalable processors and the networking ecosystem. The products are optimized for Intel® architecture and a broad operating system ecosystem: Windows*, Linux* kernel, FreeBSD*, Red Hat* Enterprise Linux (RHEL*), SUSE*, Ubuntu*, Oracle Solaris*, and VMware ESXi*. Supported connections and media types for the Intel Ethernet 700 Series are: direct-attach copper and fiber SR/LR (QSFP+, SFP+, SFP28, XLPP/CR4, 25G-CA/25G-SR/25G-LR), twisted-pair copper (1000BASE-T/10GBASE-T), backplane (XLAUI/XAUI/SFI/KR/KR4/KX/SGMII). Note that Intel is the only vendor offering the QSFP+ media type. The Intel Ethernet 700 Series supported speeds include 10GbE, 25GbE, 40GbE.

Performance results are based on testing as of the date set forth in the configurations and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit intel.com/benchmarks.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, the Intel logo, Intel Optane, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© 2019 Intel Corporation

Printed in USA

0419/RM/PRW/PDF

Please Recycle 338995-001US