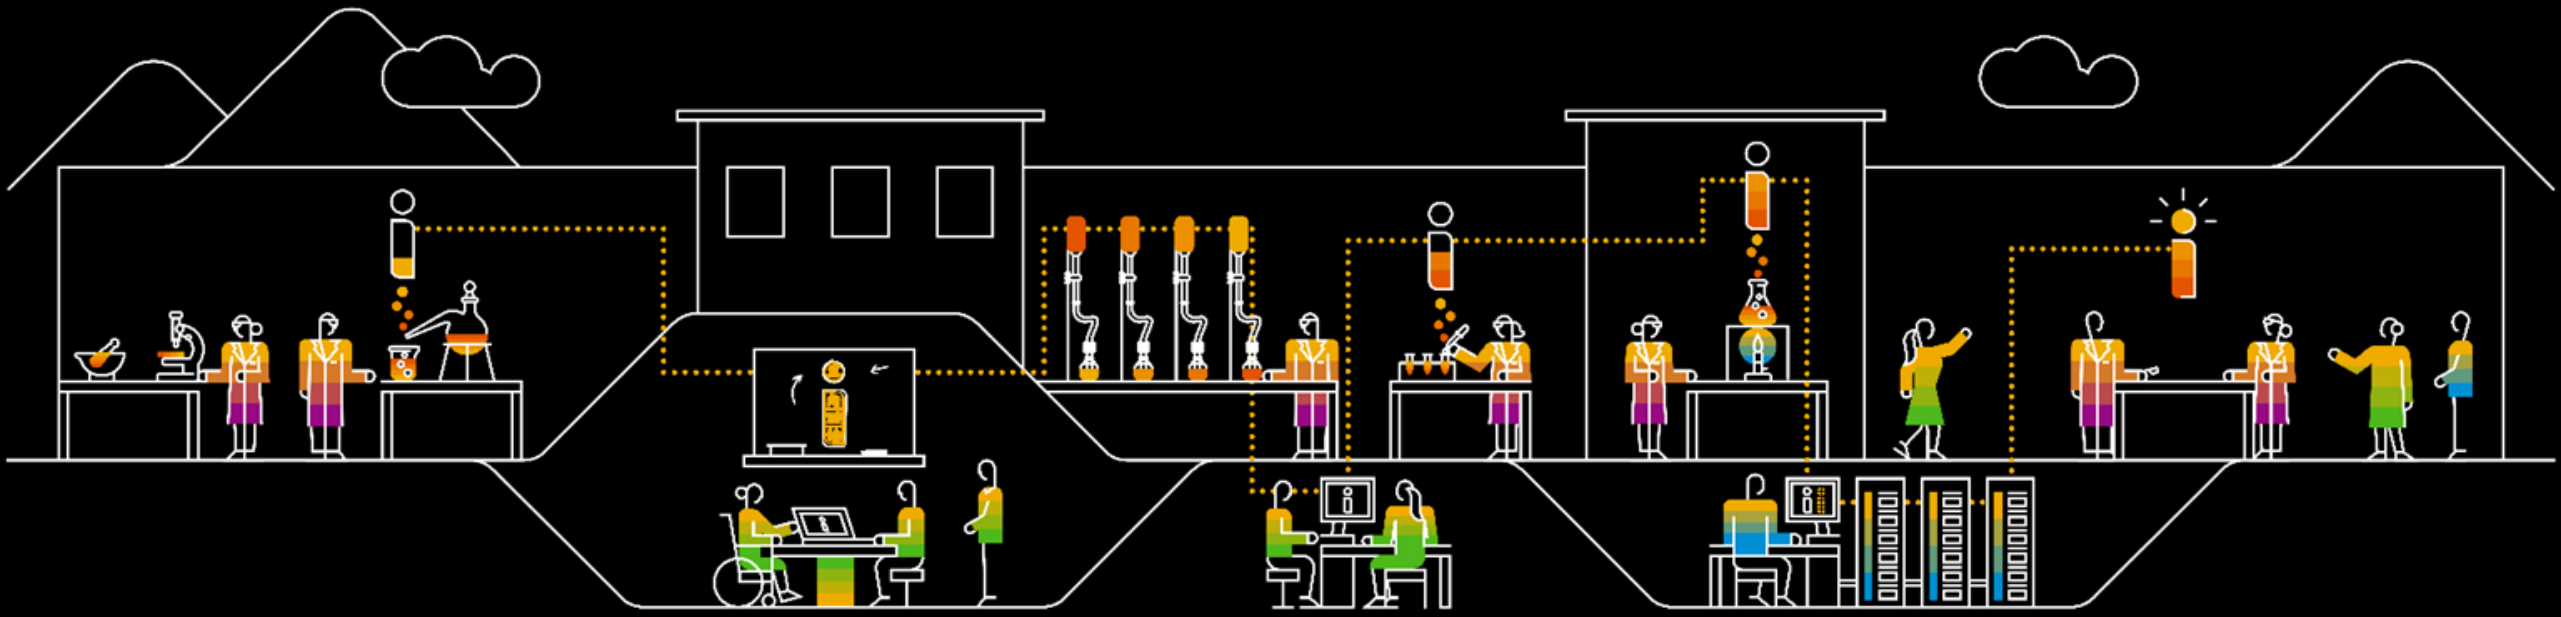

SAP Innovation Awards 2021

SAP Innovation Awards 2021 Entry Pitch Deck

Increasing scalability and performance for even the heaviest of financial workloads

PayPal Holdings Inc.

PUBLIC

THE BEST RUN SAP

Company Information

Headquarters	San Jose, California
Industry	Banking
Web site	www.paypal.com

PayPal Holdings Inc. is a leading technology platform company that enables digital and mobile payments for consumers and merchants worldwide. PayPal creates opportunity through its financial services, empowering people and businesses in new and innovative ways to help them thrive. Despite more than US$1 billion in transactions daily – including frequent reconciliation challenges – PayPal made transaction volume irrelevant to its speed of operations, ensuring its ability to grow and scale out.

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

Participating Partner Information

Google Cloud

Google Cloud offered a set of products and innovations that help maintain risk and compliance while delivering scale and reduced costs. To support the massive scale-out workload with SAP S/4HANA for financial products subledger, Google Cloud provided a number of critical capabilities used in this test configuration.

SAP S/4HANA for financial products subledger on a 96 TB SAP HANA scale-out configuration brings the best of SAP and Google Cloud engineering together to deliver groundbreaking capabilities to PayPal. By combining Google Cloud's expertise in building scalable, distributed systems with SAP's relentless focus on delivering mission-critical business process automation, PayPal now runs one of the world's largest scale-out clusters of SAP HANA, with room to scale even further.

Urs Hoelzle, SVP Technical Infrastructure, Google

Google Cloud

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

Increasing Scalability and Performance for even the heaviest of financial workloads

Paypal Holdings inc.

The SAP S/4HANA for financial products subledger solution on top of SAP HANA along with Google Cloud's services is delivering the finance system transformation we need. It is allowing us to shift from an on-prem single instance scale-up environment to a modern, multi-instance, scale-out architecture. This will allow PayPal to readily meet regional data requirements without putting additional burden on our Finance teams.

Kash Hathi, Senior Director Financial Technologies, PayPal

Challenge

Under pressure to meet regulatory requirements, deliver higher margins, and improve customer experiences, PayPal sought to manage greater transaction volumes of data with higher levels of sophistication. The limits of PayPal's current scale-up 48TB on-premise landscape did not allow for the flexibility required to meet these challenges. The cost of managing ever growing data volumes became a serious problem.

Solution

A Proof of Concept showcasing how the combination for Google Cloud and SAP S/4HANA for financial products subledger (FPSL) on a 96 TB SAP HANA scale-out configuration brings the best of SAP and Google Cloud engineering together. Providing PayPal groundbreaking capabilities with higher performance, lower cost and more flexibility using the novel native storage extension capability of SAP HANA 2.0 SPS05 on scale-out.

Outcome

Speed and robustness : Multi-node configurations in the cloud with smart data tiering allow for higher processing speeds and business-grade robustness without compromising cost or high-end performance. **Scalability:** PayPal is now able to run scale-out clusters spanning from 8 up to 10 instances, with a total of 96 TB of DRAM and 3,328 vCPUs, thus benefiting from unlimited scale and flexibility. **Performance:** Compared with a scale-up on-premise configuration, performance improved by a factor of 10 (OLTP) to 20 (OLAP), with runtime behavior linear to volume.

40x Acceleration in query runtime speed through parallelization

24 Weeks weeks of nonstop operations with zero downtime

78% reduction in memory footprint with dynamic data tiering techniques

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

Business Challenges and Objectives

Paypal had the following key challenges:

- Need to react quickly to regulatory changes on the globe and meet data processing requirements in country.
- Need to be able to scale without spending hundreds of thousands of dollars on single-purpose large scale-up machines
- Need to manage cost-effective hyper data growth without introducing complexity for operations
- Need to maintain high performance to gain timely insights from their day-end and period-end processing data

Proof of Concept objectives included:

- Showcase that the the combination of Google Cloud and FPSL on SAP HANA can scale with high performance for very large data volumes by generating 18+TB of Payments data and processing them
- Showcase that the native storage extension capability of SAP HANA 2.0 SPS05 provides an elegant solution for maintaining warm data and tier the previous fiscal year into it.
- Enable Paypal to quickly react to local data processing requirements changes by architecting a solution that is replicable and repeatable without the need of special hardware or infrastructure
- Deliver a more stable environment with massive reduction in planned maintenance time and the ability to add ressources when needed.

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

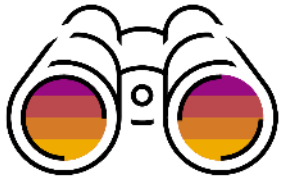

Project or Use Case Details

We focus on a payment scenario, driven by specialized payment providers. While the business logic is relatively simple, implications of data volumes and growth are not. Process optimization focuses mainly on reading and writing vast amounts of data in shorter and shorter periods. At SAP headquarters in Walldorf, Germany, engineers from Google Cloud and SAP recently performed comprehensive tests that used a total of 10 million current accounts with daily postings of 40 million business transactions (payments). The engineers processed a complete fiscal year, totaling to 12 billion transactions in Source Data Layer (SDL) and over 100 billion sub-ledger items to be managed at year end.

Half of the payment transactions are multi-currency postings (MCA), in which a payment is posted in a transaction currency that is different from the functional currency, such as when a transaction is posted in USD on a company code with leading currency EUR. The MCA transactions were the most complex in the test scenarios. MCA transactions are expensive as they need to re-evaluate the currency exchange (USD to EUR), consider the current FX spot rate, plus determine any FX losses and gains, due to FX rate changes.

The test system used 18 TB of financial data. The central subledger document tables (RDL) held more than 200 billion records, generated out of 37 billion payments. When a fiscal year is being worked on, data distribution changes significantly. At the end of the year, about 90% (16,2 TB) of the data volume is stored in RDL, making it the most sensitive data layer.

The tests provided measurement data as well as functional verification for vendors and SAP. The scope of this work was to obtain performance data, and to demonstrate scalability of SAP S/4HANA FPSL powered by SAP HANA as the database on Google Cloud.

The memory-optimized VM instance types are the preferred choice for large SAP HANA database nodes.

The first tests used the m1-ultramem-160 memory-optimized VM instances with 4 TB of DRAM and 160 vCPUs per node. The cluster initially spanned 8 nodes, with a total of 32 TB of DRAM and 1280 vCPUs. The cluster was later extended to 10 nodes to show how easy it is to scale the cluster by adding VM instances. The second and main SAP HANA scale-out landscape that is shown in the figure on slide 9 used the m2-ultramem-416 memory-optimized VM instance type that runs on the latest Intel Xeon CPU (Cascade Lake) and has 12 TB of DRAM and 416 vCPUs spread across 8 sockets. The SAP HANA scale-out cluster spanned 8 instances, with a total of 96 TB of DRAM and 3,328 vCPUs.

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

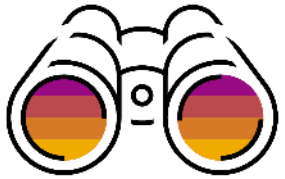

Project or Use Case Details

Results:

40x acceleration in query runtime speed through parallelization

2.7x acceleration in full-day OLTP processing

78% reduction in memory footprint with data tiering techniques

200 bn records touched in 30 seconds

10x Performance improvements for OLTP scenarios compared to scale-up architecture

20x Performance improvements for OLAP scenarios compared to scale-up architecture

24 weeks of nonstop operations with zero downtime

50% reduction in maintenance downtime window for parallelized, engineered redistribution of data

99.95% availability due to multi-regional cloud storage backups, with comprehensive encryption and no egress cost

>5GB/s consistent encrypted backup speed with native Backint integration

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

Benefits and Outcomes

Business or Social

Scalable platform to handle high volumes and hyper growth

Multi-country / entity capabilities to support global expansion or change in regulation (eg. Brexit, India) and acquisition strategy (eg. Braintree, Venmo)

Finishing each day in a day with no weekend or period end breaks

Redundancies and data movement got eliminated

Full transparency to applied finance rules at scale

IT (optional)

Business continuity with Google Live migration eliminating maintenance downtimes on infrastructure

Unparalleled performance with 20x improvement in OLAP processing & 2.7x in OLTP processing compared to scale-up systems

Massive cost savings by tiering 78% of data into native storage extension with little performance impact

Scale with more nodes when needed and adapt to data growth with highest flexibility cutting the necessary planned maintenance by 50%

Human Empowerment

Saving hours for IT personnel and business users to react to local data protection regulation changes by templating the solution and make it deployable in country

Performance improvements of the solution empowers business users to not wait for hour long processing

Predictable system behavior allows IT personnel to plan for upgrades and maintenance with less risk to business users

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

Architecture

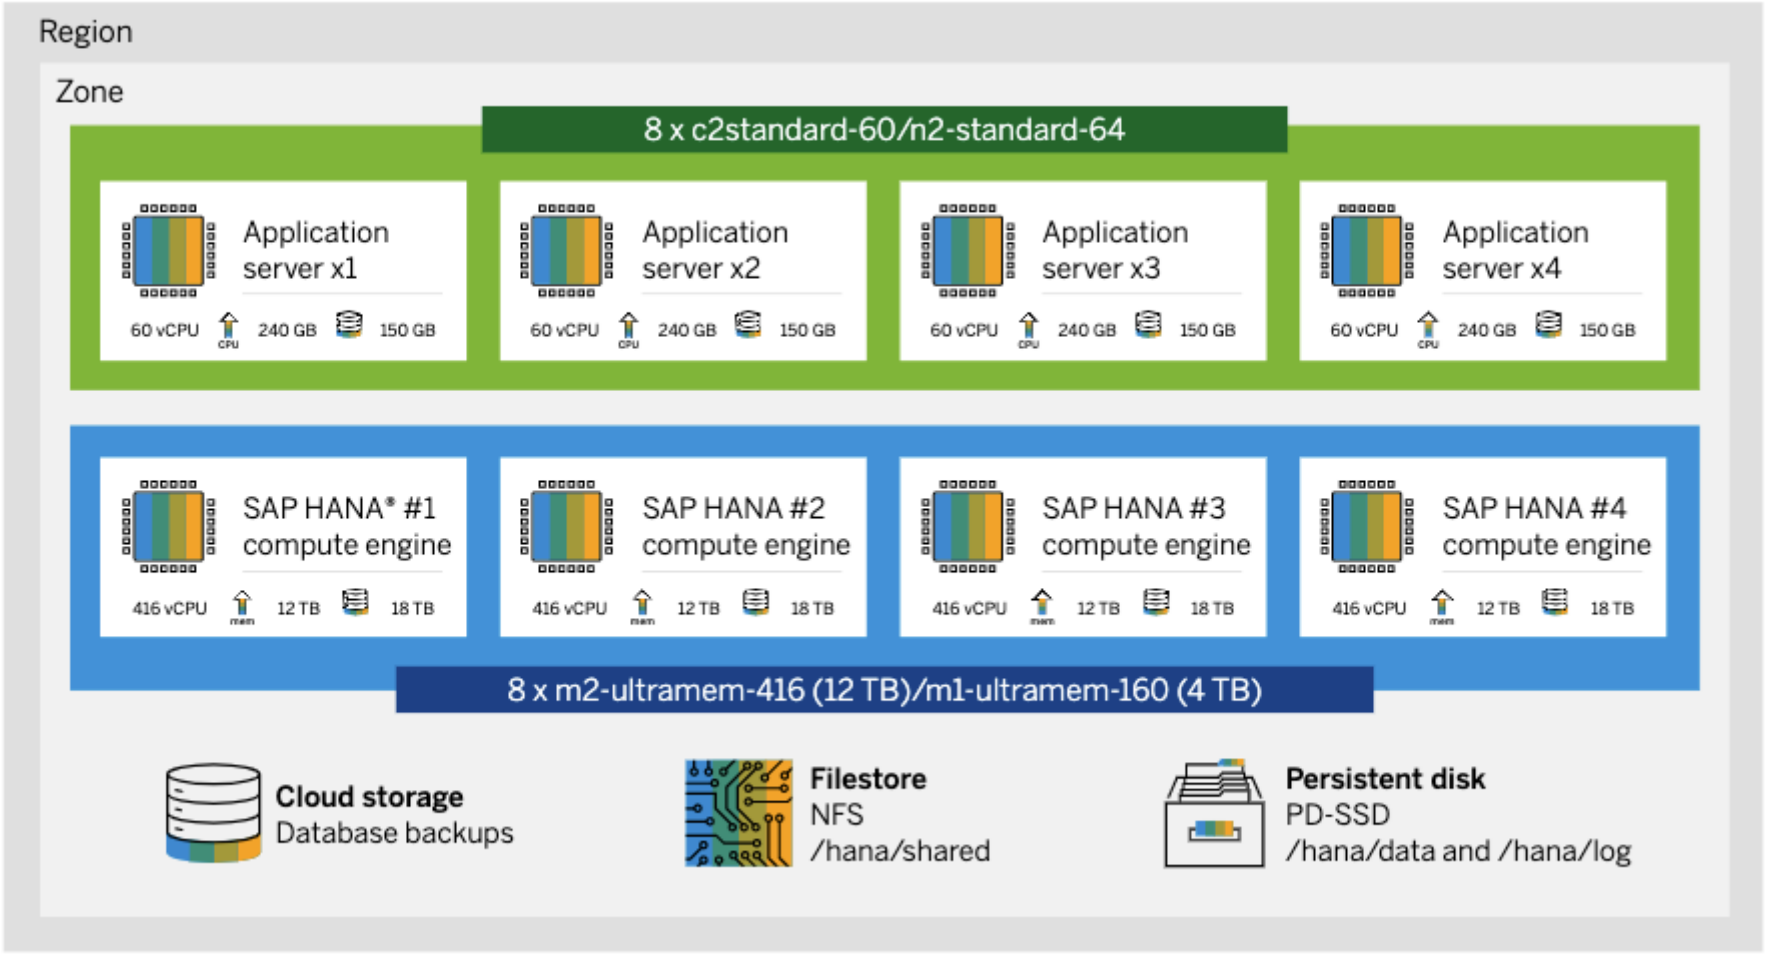

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

Deployment

Deployment status POC

Date April-September 2020

Number of users not applicable

SAP® technologies used:

	SAP product	Deployment status (live or proof of concept [POC])	Contribution to project
1	SAP S/4HANA for financial products subledger	POC	Processing payments in the cloud at high volumes in something close to real time
2	SAP HANA 2	POC	Native storage extension for scale-out configurations

If you have used one or more of the services or support offerings from SAP Services and Support during the implementation or deployment phase, please indicate which one(s) below with an X

- ☐ SAP MaxAttention™
- ☐ SAP ActiveAttention™
- ☐ SAP Advanced Deployment
- ☐ SAP Value Assurance
- ☐ SAP Model Company
- ☐ Others:
- ☐ SAP Innovation Services
- ☐ SAP Innovative Business Solutions

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

Advanced Technologies (1 of 2)

The following **advanced technologies** were part of the project.

	Technology or use case	Product used*	Contribution to project and how product used integrates with SAP products
1	**Machine learning or artificial intelligence** **Robotic process automation**, conversational AI, AI-based knowledge graph		
2	**Intelligent data management** Multi-cloud, data virtualization and governance, smart data tiering, persistent memory, data privacy	SAP HANA native storage extension	Dynamic data tiering in a scale-out environment in the cloud
3	**Advanced and augmented analytics** • Real-time and streaming analytics, spatial analytics • Natural language query and generation • AutoML to identify trends, patterns, outliers • Predictive analytics (time series analysis and forecasting, regression, classification)		
4	**Data and analytics solutions in the cloud** • Unified data and analytics cloud platforms by SAP • Modern/self-service data to analytics	SAP S/4HANA for financial products subledger (FPSL)	Processing payments in the cloud at high volumes in something close to real time

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

*If this is not an SAP product, explain how it integrates with SAP products.

Advanced Technologies (2 of 2)

The following **advanced technologies** were part of the project.

Technology or use case	Product used*	Contribution to project and how product used integrates with SAP products
5 **Advanced cloud integration** • API economy (monetization and API marketplaces) • AI-based or crowdsourced integration • High throughput, low-latency digital integration hub		
6 **Industry cloud platform**		
7 **Blockchain**		
8 **Internet of Things**		
9 **3D printing**		

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC

*If this is not an SAP product, explain how it integrates with SAP products.

Additional Information

White paper Read on sap.com / Download

Infographic: Read on google.com

Exec blog SAP: Read on sap.com

Exec blog Google Cloud: Read on google.com

© 2020 SAP SE or an SAP affiliate company. All rights reserved. I PUBLIC