



2026 High Water Benchmark Report

May 2026

Information in this document is subject to change without notice.

No part of this document may be reproduced or transmitted in any form or by any means, for any purpose, without the express written permission of TEMENOS HEADQUARTERS SA.

© 2022 Temenos Headquarters SA - all rights reserved.



Table of Contents

About this document4

 Project Resources..... 4

 List of Abbreviations and Terms Used in the Document 5

Introduction.....6

Detailed Summary.....6

 Success Criteria – Lending ,TPH ,Deposits, Microservice and copilot..... 6

High-Level Architecture7

Product Details8

 New Technologies Implemented 8

Hardware & Software Specifications9

Execution & Monitoring Tools Used 10

Scalability Test Results 11

Transaction Mix 16

Overall TPS – Transact & Microservice 19

Appendix 20

Document History

Author	Version	Date
Gayathri Parameswaran	1.0	12/05/2026

Reviewer	Version	Date
Arun Kumar Ramasamy	1.0	21/05/2026
K Bharath	1.0	21/05/2026

Comments:

About this document

This document presents a comprehensive report on the performance tests conducted for the Generic RES APIs, TPH, Deposits modularity APIs, Microservices and copilot on the latest R26 in the Azure SaaS Cloud.

Project Resources

Name	Role
Arun Kumar Ramasamy	Benchmark Lead
P Mahalakshmi	Deposits Modularity Benchmark Lead
T Mohanraj	TPH Benchmark Lead
Rakesh Karripot Balakrishnan	Architect
Dinesh Kumar Kothandaraman	Microservice and Copilot Lead
Gayathri Parameswaran	Deposits Modularity Benchmark Lead
Kiruthikkumar Sivashanmugam	Performance Engineer
Prakash V	Performance Engineer
Praveen Lalbahathur	Performance Engineer
Aiswarya Somi Gopinath	Performance Engineer
Someshvar V	Performance Engineer
Koustav Das	Performance Engineer
Syed Amanat Abbas	Performance Engineer
Poobalan Kumar	Database Administrator

List of Abbreviations and Terms Used in the Document

Abbreviation/Term	Expansion/Description
HWBM	High Water Benchmark
ACA	Azure Container App
MS	Micro-Services
API	Application Programming Interface
TPS	Transaction Per Second
ms	milliseconds
ALT	Azure Load Testing
SEAT	System Engineering Automation Tool
AKS	Azure Kubernetes Service
RES	Retail Enterprise Services
SLA	Service Level Agreement
OLB	Online Banking
GAI	Generic Accounting Interface
ARM	Advanced RISC Machine
SKU	Stock Keeping Unit
TPH	Temenos Payment Hub
PO	Payment Order
TDH	Temenos Data Hub

Introduction

The 2026 Highwater Benchmark was conducted to evaluate the performance, scalability, and sustainability of Temenos product offerings on the Microsoft Azure Platform. The benchmark demonstrates improvements achieved compared to the previous year and highlights the adoption of innovative technologies across various Temenos solutions and microservices. Performance testing was carried out for Holdings MS, Party Microservices, Temenos Payments Hub (TPH), and Deposits modularity to validate their capability to handle large-scale workloads efficiently in a cloud-native environment.

Executive Summary

- Overall, **16,385** TPS were achieved.
 - 1,539 TPS on Lending Transactions
 - 244 PPS on TPH
 - 2,109 TPS on Deposits
 - 12,241 TPS on Enquiries (Holding Microservices).
 - 252 TPS on Copilot for core

Detailed Summary

Success Criteria – Lending, TPH, Deposits, Microservice and Copilot for Core

In this exercise, a transaction mix of 76.2% Enquiries and 23.8% Transactions was used by targeting 10K TPS, with the detailed breakdown provided in the table below.

API	Component	Achieved TPS	Response Time(ms)
Customer Creation	Lending	80	123
Current Account Creation	Deposit	25	662
Savings Account Creation	Deposit	31	549
Savings Account Creation	Lending	55	945
Loan Creation	Lending	29	1,810
Deposit Creation	Deposit	28	604
Deposit Funding	Deposit	35	485
Deposit – Credit Booking	Deposit	1,006	90
Deposit – Debit Booking	Deposit	984	92
Customer Enquiry	Party	1,080	10
Balance Enquiry	Holdings	4,845	5
Transaction List	Holdings	2,988	6
Loan Information	Holdings	1,584	8
Deposit Information	Holdings	1,744	7
Incoming PO	TPH	129	213
Outgoing PO	TPH	115	412
Loan Repayment	Lending	215	160
ATM – DEBIT+CREDIT	Lending	1,160	105
Copilot for core	Copilot	252	23
Total		16,385	

Dataset:**Transact:**

- 50 million Customers, 25 million Savings Accounts, 25 million Current Accounts, and 25 million Loan Accounts.
- A total of 6 transactions were loaded for each Savings and Current Account.
- All the loans were created and disbursed to customers' savings accounts.


Deposits:

- 25 million Deposit Accounts, 25 million Current Accounts.

Microservices:

- 50 million Customers, 125 million Accounts

High-Level Architecture



Product Details

Lending:

Retail Banking Enterprise Services (RES) is a ready-to-deploy solution that covers the full range of retail banking operations, combining Temenos Infinity, Transact, and Financial Crime Mitigation (FCM). It includes preconfigured products for accounts, lending, deposits, and multi-currency services, helping banks accelerate innovation and respond quickly to market needs. The solution is deployed on AKS, integrated with an Azure SQL database.

Deposits Modularity

A standalone module that provides a specific set of functions and features, which is independently upgradable and can operate as a point-solution or form part of a larger landscape with other Temenos or third-party solutions.

Temenos Payments Hub (TPH):

Temenos Payments Hub Messaging Framework is a middleware that is used to accept and generate messages. It allows the exchange of messages between Temenos Payments Hub Standard/Temenos Payments Hub Enterprise (TP/TPH) and other external systems. The solution is deployed on AKS, integrated with an Azure SQL database.

Holdings & Party Microservices

Holdings Microservices is the solution for holding the accounting entries & customer account balances, and this microservice is used to perform enquiries on account balances, account transactions list, loan balances, loan transactions list and loan details as part of this benchmark. Party Microservices is the solution for managing customer related information like customer personal information, address, etc.

Both microservices were deployed in the Azure Container Apps, integrated with Azure PostgreSQL as a database.

Copilot for Core

Copilot for Core is an AI-powered assistant embedded in Temenos Core. It is designed to improve productivity and help the bank users to retrieve customer information instantly, monitor branch performance, and identify potential sales opportunities with less effort.

With Copilot for Core, users can:

- Ask natural-language questions about Temenos Core data.
- Fetch real-time information from Temenos Core.
- Accelerate daily operations by reducing manual navigation.

New Technologies Implemented

Technologies	Summary
ARM and ARM Cobalt	ARM and ARM Cobalt refer to Microsoft's underlying architecture for the SKU, offering better performance compared to the Intel architecture.
JDK & Wildfly	The version has been upgraded from wf26jdk11 to wf32jdk21 for base image
Jakarta	To run the Java Platform, Jakarta has been replaced with Javax.
One Container & Direct Connect	This technology is set up within a single container, where both the IRF-Provider and the Browser Web are hosted.

Hardware & Software Specifications

Hardware Specification

Component	Node Configuration	Pod Configuration	
		vCPU	Memory (GB)
Lending and TPH	Standard D96plds v6 (96 vCPU, 192 GiB memory)	6	12
Deposits	Standard D96plds v6 (96 vCPU, 192 GiB memory)	6	13
Holdings MS	General Purpose Dedicated D-32(32 vCPU, 128GB memory)	1	2
Party MS	General Purpose Dedicated D-16(16 vCPU, 64GB memory)	1	2
Copilot	Standard_D16ds_v5 (16 vCPU, 64 GB memory)	6	12
TDH	Standard D96lds v5 (96 vCPU, 192 GiB memory)	4	10

Database Layer

Application	Database	Compute		Storage	vCPU	Memory (GB)
		Tier	Size			
Lending and TPH	Azure SQL Database	Hyperscale	Premium series	Hyperscale SSD (1.95TB)	128	625
Deposits	Azure SQL Database	Hyperscale	Premium series	Hyperscale SSD	192	995
Holdings MS	Azure PostgreSQL	General Purpose	Standard_D96ds_v5	Premium SSD (2048 GB)	96	384
Party MS	Azure PostgreSQL	General Purpose	Standard_D16ds_v5	Premium SSD (512 GB)	16	64
Copilot – TBAC PostgreSQL	Azure PostgreSQL	General Purpose	Standard_D4ds_v5	32 GB	4	16
Copilot – ODS Database	Azure SQL Database	Hyperscale Gen5	Standard-series (Gen5)	961 MB	2	10.4
TDH	Azure SQL Database	Hyperscale PRMS	Premium series	Hyperscale SSD 1TB	80	560

Execution & Monitoring Tools Used






- The Load for Transact, Deposits, Copilot and Microservice was executed via ALT.
- The Load for Embedded TPH was done via J-Meter.
- The following tools were used to monitor the benchmark tests

Monitoring Tool	Purpose
Dynatrace	To review the resource utilization, application performance monitoring, in-depth analysis.
Grafana	To review the resource utilization, NGINX, Ingress performance metrics
Azure Database Watcher	Enabled to monitor Azure databases, providing real-time insights, performance metrics, and diagnostic logs to detect issues and optimize queries automatically.
JMeter	For injecting the required data volume and measuring scalability.
Azure Load Testing (ALT)	For executing the API test runs to capture the scalability results and the targeted metrics
TAFJ Session Monitor	Enabled during the runs to detect the locking/sleep and to verify active sessions/agents during the execution.
Session Management	Enabled during the runs to detect the heap memory of the application.
SEAT	SEAT is the core tool for transaction analysis, capturing key metrics like reads, writes, I/Os, path length, and elapsed time. It helps evaluate transaction performance and system efficiency.
Queue Metrics Tool	Custom tool developed using Python language to visualize the ActiveMQ-level metrics – TPS, Response Time as separate graphs


Scalability Test Results

Lending Vertical Scalability


Expected TPS Achieved TPS





Lending Horizontal Scalability



TPH Vertical Scalability







TPH Horizontal Scalability



Deposit Vertical Scalability

Expected TPS Achieved TPS





Vertical Scalability Graph





Deposit Horizontal Scalability


Expected TPS Achieved TPS







Microservices Scalability

Expected TPS Achieved TPS


Party MS



Holdings MS



Copilot Scalability



Transaction Mix

For transaction mix, the load test was conducted for below list of APIs with a 30-minute ramp-up time followed by a 30-minute stable period.

Lending – Test Metrics:

Test Cases	Pods	Total Vcpu	Achieved TPS	TPS / Threads	Response Time (ms)	
					Avg(ms)	p95(ms)
POST_LoanCreation	9	54	29	0.54	1810	2900
POST_SavingsAccount	10	60	55	1.06	945	1700
POST_LoanRepayment	7	42	215	6.14	160	252
POST_CsmBook_Debit	6	36	660	11	90	110
POST_CustomerCreation	1	6	80	8	123	142
POST_CsmBook_Credit	6	36	500	8.33	120	198

Lending - Resource Utilisation

Test Cases	App (%)		Database (%)	
	CPU%	Memory%	CPU%	Memory%
	(max)	(avg)	(max)	(avg)
POST_LoanCreation	86%	97%	11%	28%
POST_SavingsAccount	85%	96%		
POST_LoanRepayment	48%	97%		
POST_CsmBook_Debit	81%	98%		
POST_CustomerCreation	90%	97%		
POST_CsmBook_Credit	95%	97%		

TPH – Test Metrics:

Test Cases	Pods	Threads / Pod	Total vCPU	Achieved TPS	TPS / Theads	Response Time (ms)	
						Avg(ms)	p95(ms)
Incoming Payments	4	8	24	129	4.0	213	237
Outgoing Payments	6	8	36	115	2.4	412	630

TPH - Resource Utilisation

Test Cases	App (%)		Database (%)	
	CPU% (avg)	Memory% (avg)	CPU% (avg)	Memory % (avg)
Incoming Payments	73.3	90	11	29.4
Outgoing Payments	80.3	90		

Deposits Mix Test Results:

Apis	Pods	Threads / Pod	Total vCPU	Achieved TPS	TPS / Threads	Response Time (ms)	
						Avg(ms)	p95(ms)
Credit Book	16	6	96	1006	10.5	90	112
Debit Book	16	6	96	984	10.3	92	114
Payment Funding	3	6	18	35	1.9	485	597
Savings Account	3	6	18	31	1.7	549	633
Term Deposits	3	6	18	28	1.6	604	735
Current Account	3	6	18	25	1.4	662	799

Resource Utilisation

Apis	App (%)		Database (%)	
	CPU% (avg)	Memory% (avg)	CPU% (avg)	Memory% (avg)
Credit Book	68%	83%	10%	43%
Debit Book	70%	84%		
Payment Funding	75%	79%		
Savings Account	79%	81%		
Term deposits	78%	81%		
Current Account	79%	82%		

Microservice – Holdings – Test Metrics:

Use Cases	TPS	TPS/ Pod	TPS/	Avg (ms)	p95(ms)	Replicas
			Thread			
GET_BalanceEnquiry	4,845	385	285	3	6	29
GET_TransactionList	2,988		249	3	6	
GET_BulkArrangementLoan	1,584		176	5	8	
GET_BulkArrangementDeposit	1,744		218	4	7	

Party – Test Metrics:

Use Cases	TPS	TPS/ Pod	TPS/	Avg (ms)	p95(ms)
			Thread		
GET_PartyCustomer	1,080	1,080	135	7	10

Microservice - Resource Utilisation**Holdings**

Use Cases	Pod Metrics		Database Metrics	
	CPU% (avg)	Memory% (avg)	CPU% (avg)	Memory% (avg)
GET_BalanceEnquiry	92	36	3	28
GET_TransactionList				
GET_BulkArrangementLoan				
GET_BulkArrangementDeposit				

Party

Use Cases	Pod Metrics		Database Metrics	
	CPU% (avg)	Memory% (avg)	CPU% (avg)	Memory% (avg)
GET_BalanceEnquiry	95	25	6	33

Copilot – Test Metrics:

Test Cases	Pod	Total Vcpu	Threads / Pod	Achieved TPS	TPS / Threads	TPS/ vCore	Response Time (ms)	
							Avg(ms)	p95(ms)
GET_Conversation	1	3	2	135	67.7	45	13.5	31
GET_Conversation_History			2	117	58.5	39	32.2	41.3

Copilot - Resource Utilisation

Use Cases	Pod Metrics		Azure Postgress Database Metrics	
	CPU% (avg)	Memory% (avg)	CPU% (avg)	Memory% (avg)
GET_Conversation	76	54.9	11.9	35.8
GET_Conversation_History				

Overall TPS – Lending, Deposits, TPH & Microservice



Conclusion

- The high-water benchmarking for TCF 2026 was completed successfully, achieving more than the targeted 16,385 TPS across Transact, TPH, Deposits, Microservices, and Copilot.
- For Deposits Modularity, a total throughput of 2,109 TPS was achieved, exceeding the target TPS with optimal application-side resource utilization. All APIs demonstrated near-linear scalability in both vertical and horizontal scaling scenarios and successfully met the defined SLA requirements for response times.
- For Microservices, compared to last year, a new Deposit Enquiry API was introduced under Bulk Arrangement for Holdings, and the Customer Enquiry API migrated from Transact to Party, resulting in an 82% improvement in response time. Additionally, database resource utilization for Holdings was reduced by 57% during the Mix Run.
- Copilot for Core – This year, we have introduced Temenos AI powered is introduced in the benchmark transaction mix. In Copilot for Core, NLQ Agents covering prompts querying ODS database and Conversation APIs were tested, and target throughput was successfully achieved as part of the benchmark execution.

Appendix

- Confluence Page

[HWBM-2026 - Technology / KTLO - Assurance - Confluence PROD](#)