Interaction Concentrator 8.1

Sizing Guide Whitepaper

ICON 8.1.5 Performance Testing

About Genesys

Genesys is the world's leading provider of customer service and contact software - with more than 4,000 customers in 80 countries. Drawing on its more than 20 years of customer service innovation and experience, Genesys is uniquely positioned to help companies bring their people, insights and customer channels together to effectively drive today's customer conversation. Genesys software directs more than 100 million interactions every day, maximizing the value of customer engagement and differentiating the experience by driving personalization and multi-channel customer service - and extending customer service across the enterprise to optimize processes and the performance of customer-facing employees. Go to www.genesys.com for more information.

Each product has its own documentation for online viewing at the Genesys Technical Support website or on the Documentation Library DVD, which is available from Genesys upon request. For more information, contact your sales representative.

Notice

Although reasonable effort is made to ensure that the information in this document is complete and accurate at the time of release, Genesys Cloud Services, Inc. cannot assume responsibility for any existing errors. Changes and/or corrections to the information contained in this document may be incorporated in future versions.

Your Responsibility for Your System's Security

You are responsible for the security of your system. Product administration to prevent unauthorized use is your responsibility. Your system administrator should read all documents provided with this product to fully understand the features available that reduce your risk of incurring charges for unlicensed use of Genesys products.

Trademarks

Genesys and the Genesys logo are registered trademarks of Genesys Cloud Services, Inc. All other company names and logos may be trademarks or registered trademarks of their respective holders. © 2012 Genesys Cloud Services, Inc. All rights reserved.

The Crystal monospace font is used by permission of Software Renovation Corporation, www.SoftwareRenovation.com.

Technical Support from VARs

If you have purchased support from a value-added reseller (VAR), please contact the VAR for technical support.

Ordering and Licensing Information

Complete information on ordering and licensing Genesys products can be found in the ***Genesys Licensing Guide***.

Released by: **Genesys Cloud Services, Inc.** www.genesys.com

Document Version: 815icon_sizing_supp_03-2022

Table of Contents

Table of Figures

GENESYS

Preface

Welcome to the *Interaction Concentrator 8.1 Sizing Guide Whitepaper.* This document is a supplement to the Interaction Concentrator chapter in the *Genesys Hardware Sizing Guide*.

This document is valid only for the 8.1.5 releases of this product.

For information about related resources, see the supplementary material starting here.

About Interaction Concentrator

Interaction Concentrator collects and stores detailed data about the interactions and resources in customer interaction networks that use Genesys Framework (contact center, enterprise-wide, or multi-enterprise telephony and computer networks).

Interaction Concentrator consists of two components:

- Interaction Concentrator (ICON) server
- Interaction Database (IDB)

Note: Before proceeding, review the product overview chapter in the *Interaction Concentrator Deployment Guide*, to familiarize yourself with the product architecture.

Note: Because of their variability and complexity, large contact centers require special planning. If you want to plan an Interaction Concentrator deployment for a large contact center, contact Genesys Professional Services.

Related Documentation Resources

The following resources provide additional information that is relevant to this software. Consult these additional resources as necessary.

- The *Management Framework 9.x Deployment Guide*, which will help when using Configuration Manager.

Genesys

Consult these additional resources as necessary:

- The *Genesys Technical Publications Glossary,* which ships on the Genesys Documentation Library CD and which provides a comprehensive list of the Genesys and CTI terminology and acronyms used in this document.
- The Release Notes and Product Advisories for this product, which are available on the Genesys Technical Support website at https://genesyspartner.force.com/customercare/GenesysCommunityLogin.

Genesys product documentation is available on the:

- Genesys Technical Support website at https://genesyspartner.force.com/customercare/GenesysCommunityLogin.

GENESYS

Chapter

1 Operating Environment

This document describes the factors that affect Interaction Concentrator performance and lists sample performance measurements for Amazon Cloud platform for both Windows and UNIX. This document also describes testing that was conducted in a largescale environment to examine the performance and scalability of ICON 8.1.5.

This chapter contains the following sections:

Hardware Architectures

The Interaction Concentrator architecture is flexible enough to store reporting data for a contact center environment of practically any size. The hardware architecture that you select for your Interaction Concentrator deployment depends primarily on:

- The size of your contact center, in terms of the number of daily interactions and the number of agents. This determines the requirements for your Genesys Framework Configuration Layer and Media Layer components, which are the major sources of data for Interaction Concentrator.
- The type of data that you need to collect. In a multi-site environment with multiple IDBs, considerations include whether and how your downstream reporting application will merge inter-site interactions.

  You can balance the database-writing load by distributing some ICON roles among ICON instances. For more information, see the information about ICON roles in the product overview section in the *Interaction Concentrator Deployment Guide*.

For more information about the supported Interaction Concentrator deployments, see the information about deployment scenarios in the product overview section in the *Deployment Guide*.

Hardware Guidelines

The memory requirements for the ICON application depend on the specific deployment and contact center characteristics.

The following factors affect the amount of memory and processing capacity that ICON requires:

- The number of interactions, which is a function of both the overall contact center size and the complexity of interaction flows
- The ICON configuration—for example, the roles that the ICON application has been configured to perform, or configuration settings that affect persistent queue, operational memory, and database-writing operations

Because requirements are so deployment-specific, it is not possible to provide precise hardware specifications for the ICON server. However, Genesys has conducted performance tests for Interaction Concentrator on reference platforms for both Windows and UNIX. The test hardware, in combination with the reported usage results, can serve as a baseline that you can use when you calculate your own requirements.

Database Configuration and Optimization

Tablespace configuration and database tuning can significantly affect the performance of Interaction Concentrator and your downstream reporting application.

Reducing I/O Contention

Any system configuration or database strategies to optimize database input and output (I/O) will significantly improve Interaction Concentrator performance. In particular, storing RDBMS logs, indexes, and table data on different disk drives reduces I/O contention. Genesys strongly recommends that you equip your RDBMS host with a fiber array or with a disk subsystem that contains multiple SCSI disk drives.

Recommendations for High Call Volumes

For environments with high call volumes, Genesys strongly recommends a multi-spindle disk subsystem, preferably with an advanced controller with write-back cache. Genesys also recommends locating the database log and

temporary tablespace on disks that are separate from the disks where tables and indexes are stored.

In addition, for high call volumes, Genesys recommends that you configure IDB with a large buffer cache (hundreds of megabytes, if not gigabytes).

ICON Performance Testing Resources

Rationale

The primary factors that affect Interaction Concentrator performance are the following:

- IDB size
- RDBMS type
- RDBMS settings
- IDB tuning
- The speed of the network connections between components
- The amount of interactions and business data attached to interactions

Note: Unless specified otherwise, call rates that are cited in the descriptions of the test conditions are for the contact center as a whole, not for each T-Server or Interaction Server.

Genesys performance testing of ICON focused on validating the following performance-based requirements of ICON in a large-scale testing environment:

- Ability to handle a large number of active interactions in a Voice and a Multimedia environment
- Ability to catch up on a reasonable backlog without data delays

Hardware and Software

All the tests were run in an Amazon EC2 environment.

Identical tests were conducted on AWS instances for the following platforms:

- Microsoft Windows Server 2019 Datacenter (64-bit)
- Amazon Linux 2 (64-bit)

Two separate test environments were utilized—a voice-specific deployment and a multimedia (chat) deployment.

The ICON application, RDBMS server, and IDB were hosted on different machines.

Application	Host specification	Version
ICON (Windows)	OS: Microsoft Windows Server 2019 Datacenter (64-bit) CPU: 4 up to 3.0 GHz Intel Scalable Processor (Intel AVX, Intel Turbo) RAM: 16 GB Hdd: 70 GB	ICON 8.1.514.47
ICON (Linux)	OS: Amazon Linux 2, kernel version 4.14.186 (64-bit) CPU: 4 up to 3.0 GHz Intel Scalable Processor (Intel AVX, Intel Turbo) RAM: 16 GB Hdd: 30 GB	ICON 8.1.514.47
RDBMS for IDB	CPU: 2 CPU RAM: 8 GB Cloud PostgreSQL server	PostgreSQL 9.6.20
DB Server	OS: Amazon Linux 2, kernel version 4.14.186 (64-bit) CPU: 2 Intel(R) Xeon(R) Platinum 8259CL CPU @ 2.50GHz RAM: 4 GB Hdd: 30 GB	DB Server 8.1.302.07

Note: For more information about the specifications for types of Amazon EC2 instances, see the Amazon EC2 documentation.

Hardware parameters measurements

During the Genesys performance testing, respective metrics were collected, stored, and displayed using Telegraf service, InfluxDB, and Grafana.

GENESYS

Chapter

2 Voice Tests

This chapter describes the Interaction Concentrator 8.1.5 Voice testing. The chapter contains the following sections:

Voice Test Environment

The voice test environment was organized as one tenant containing five Switches and five corresponding SIP Servers. All the SIP Servers were served by a single ICON and IDB.

On each switch, the following were configured:

- 10,000 Agent Logins
- 10,000 Extensions
- 1,000 Routing Points

A total of 10,000 agents were logged in.

Call configuration included attached data with 50, 25, or 10 key-value pairs (KVPs) as described under Load specification for Voice.

Tests were conducted independently on Windows and Linux. A single ICON monitored all call activity and stored data in a single IDB. ICON options to filter the data that is stored in IDB were turned off.

Each test started with an empty IDB and empty persistent queue file (empty backlog). IDB was not purged during the test.

Two database connection scenarios were used:

- Database connection is stable during the test (empty backlog)

- Database connection is lost during the test (non-empty backlog):
  - For the first hour of the test and under a normal workload, loss of database connection is simulated. The persistent queue file grows for the hour that the database connection is broken.
  - Then database connection is restored.

Load specification for Voice

Call parameters for ICON load:

- 16.6 calls per second (cps)
- 5 SIP Servers
- 10,000 Agent login sessions

Userdata parameters:

- 50, 25, or 10 keys
- Length of key names: 32 bytes
- Length of key values: 124 bytes

The running load interval was 6 hours. Calls from 5 SIP Servers were processed in parallel. Approximately 360,000 calls in total were processed within the interval.

Call lifespan was set randomly from 60 to 600 seconds. The lifespan of login sessions ranged from 1,800 to 10,800 seconds (30 minutes to 3 hours).

The following table shows the distribution of call scenarios.

Call scenario	% of total calls
SIP Inbound Call with Treatment	80%
SIP Inbound Call and Consultation Call	4%
SIP Inbound Call with Single-Step Transfer	4%
SIP Inbound Call with Two-Step Transfer	4%
SIP Inbound Call with Single-Step Conference	4%
SIP Inbound Call with Two-Step Conference	4%

Voice Performance Test Results

Results Summary

The following table summarizes the voice performance test results for the baseline and backlog tests.

Load specification	Application	Average CPU	Max CPU	Average RAM	Max RAM
Voice Baseline performance tests • 10,000 login sessions • Call rate of 16.6 cps					
• 50 userdata keys	ICON (Windows)	15.5%	49.8%	1.07 GB	1.10 GB
	ICON (Linux)	14.8%	44.0%	0.986 GB	1.01 GB
• 25 userdata keys	ICON (Windows)	12.5%	51.5%	871 MB	886 MB
	ICON (Linux)	11.1%	78.9%	877 MB	944 MB
• 10 userdata keys	ICON (Windows)	7.87%	48.3%	758 MB	772 MB
	ICON (Linux)	9.17%	32.3%	787 MB	801 MB
Voice Baseline performance tests with backlog • 10,000 login sessions • Call rate of 16.6 cps • For the first hour ICON DB Servers (primary and backup) are stopped to prevent ICON from storing data in the database					
• 50 userdata keys	ICON (Windows)	16.8%	54.6%	1.06 GB	1.1 GB
	ICON (Linux)	9.17%	32.3%	787 MB	801 MB

Grafana Plots for Voice Performance Tests

This section shows Voice performance test data visualized through Grafana for Windows and Linux, respectively.

- Voice Performance Tests on Windows
- Voice Performance Tests on Linux

Voice Performance Tests on Windows

The screenshots in this section show Grafana graphs and metrics for ICON 8.1.5 Voice baseline tests on Windows with 50 userdata keys, 25 userdata keys, 10 userdata keys, and a backlog, respectively.

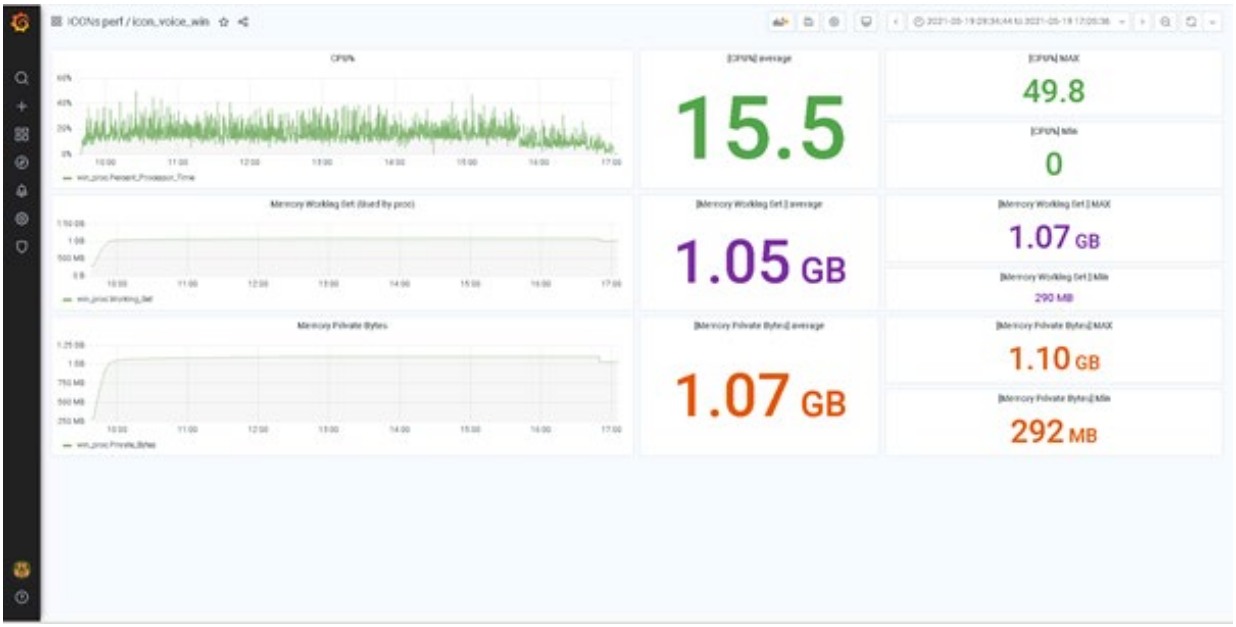

Figure 1: Voice baseline performance test — 50 userdata keys — Windows

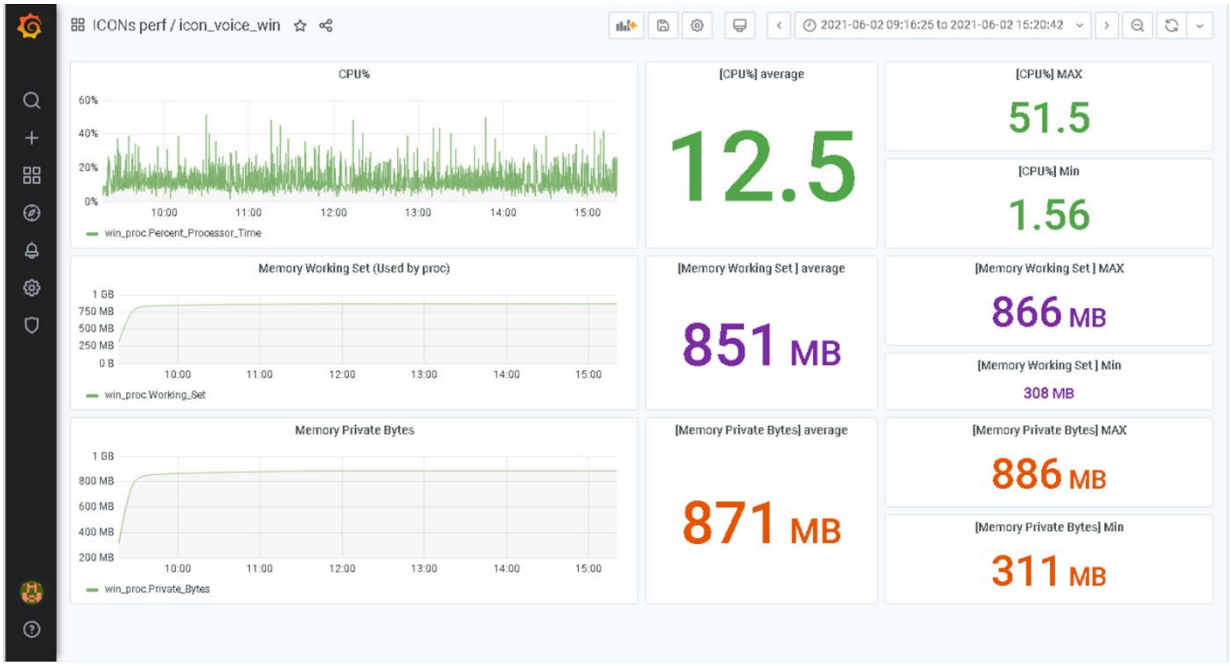

Figure 2: Voice baseline performance test — 25 userdata keys — Windows

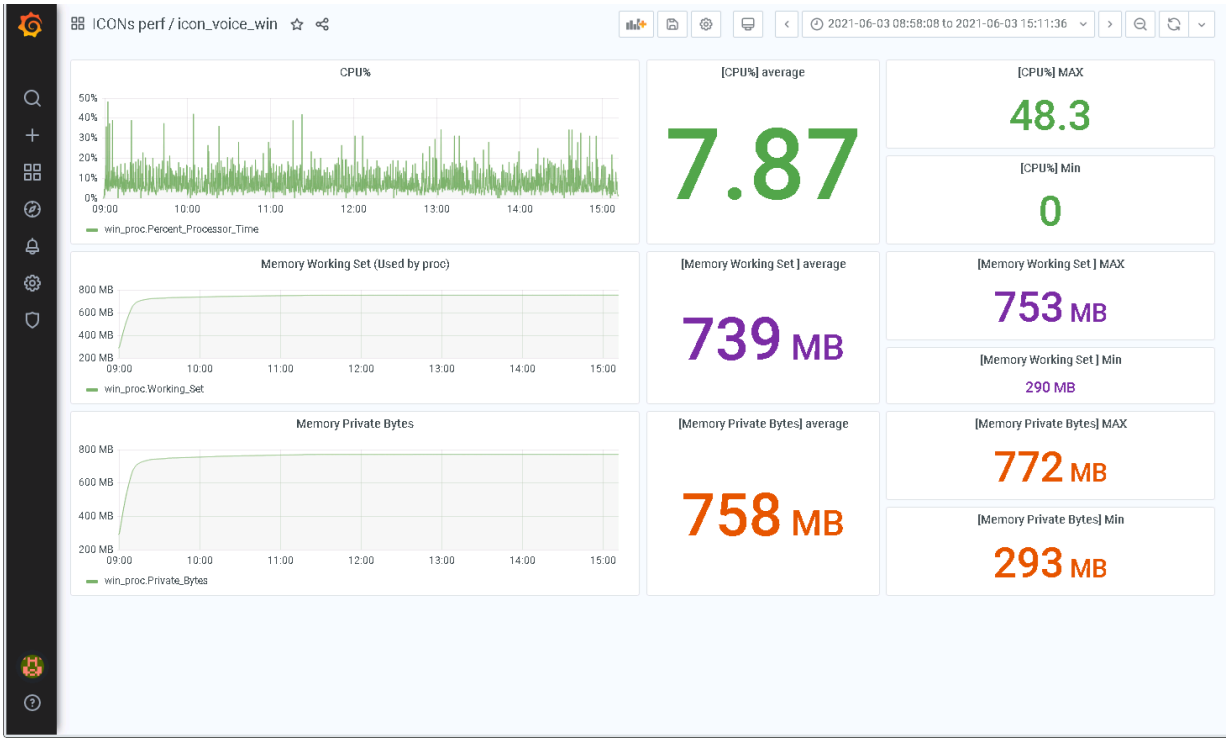

Figure 3: Voice baseline performance test — 10 userdata keys — Windows

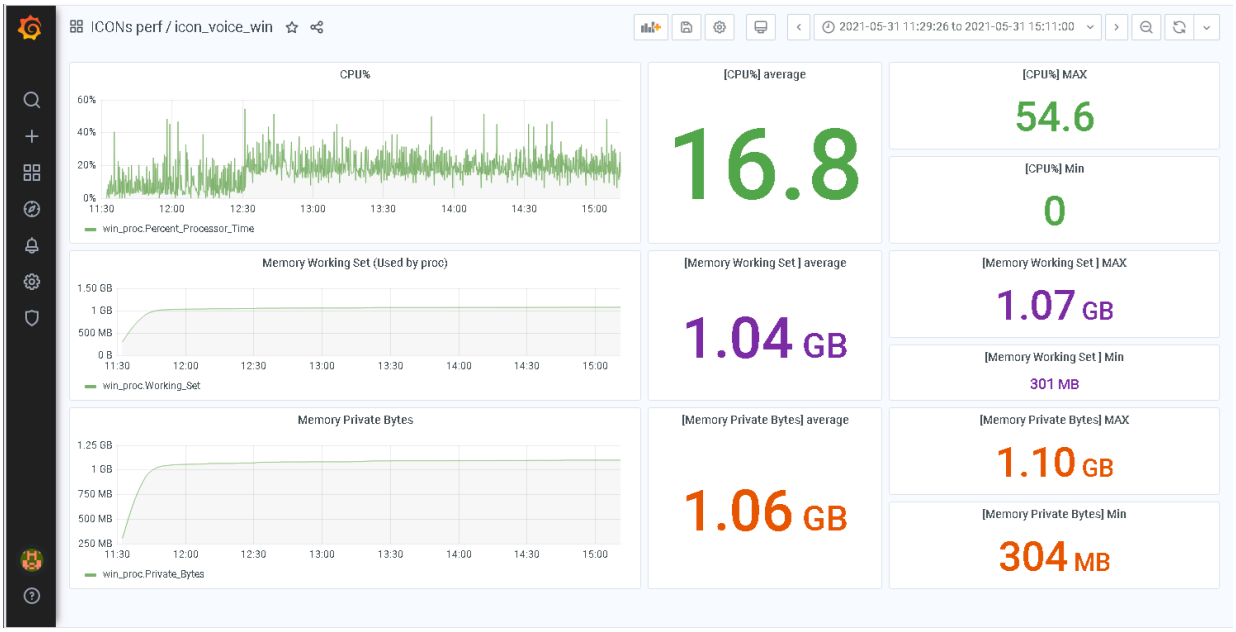

Figure 4: Voice baseline performance test with backlog — Windows

Voice Performance Tests on Linux

The screenshots in this section show Grafana graphs and metrics for ICON 8.1.5 Voice baseline tests on Linux with 50 userdata keys, 25 userdata keys, 10 userdata keys, and a backlog, respectively.

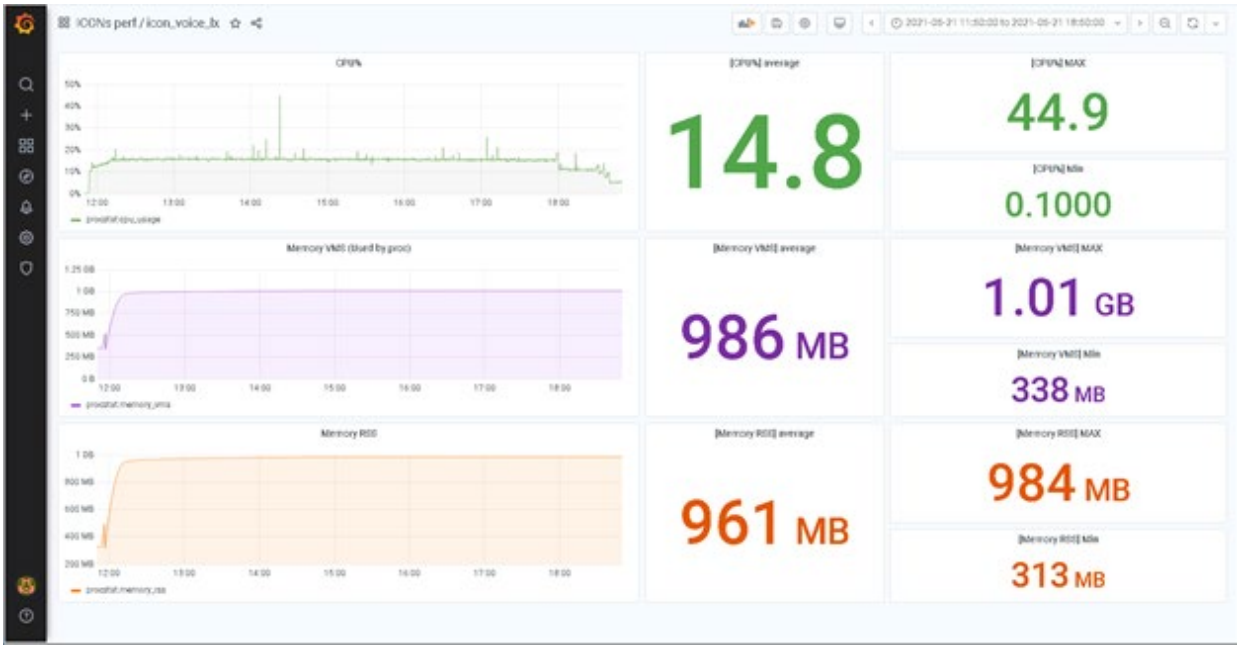

Figure 5: Voice baseline performance test — 50 userdata keys — Linux

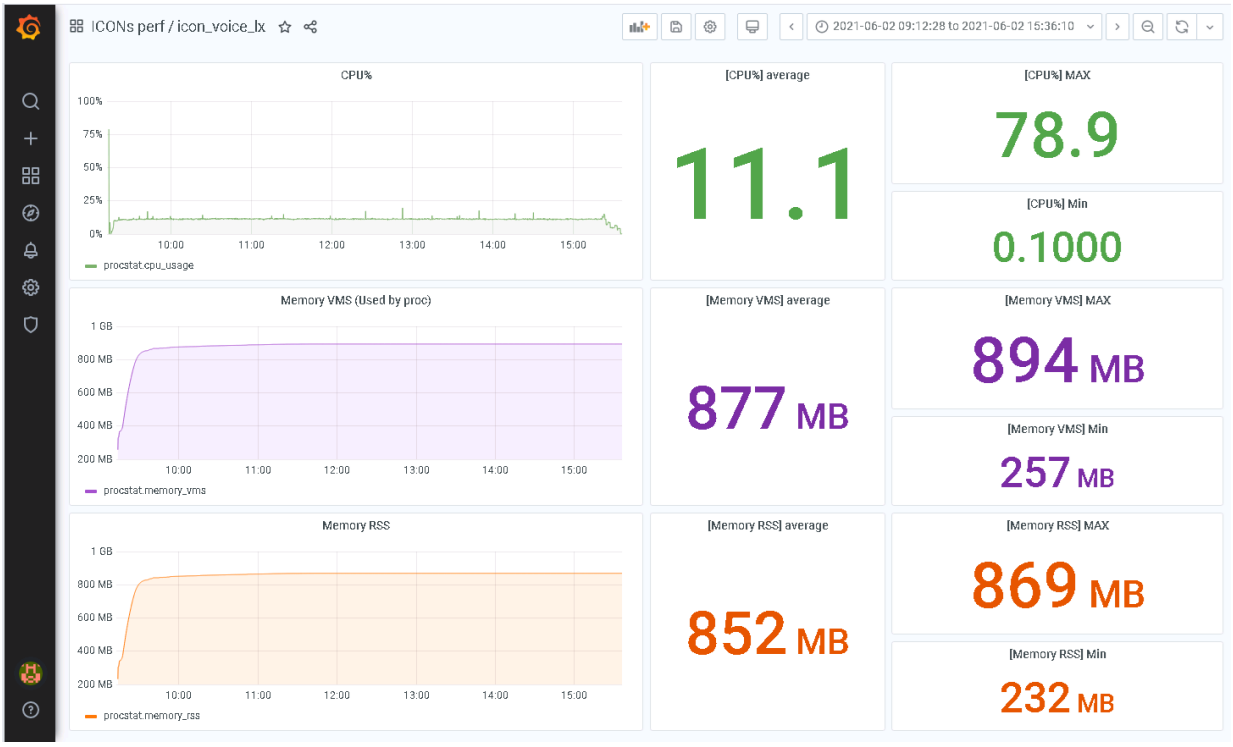

Figure 6: Voice baseline performance test — 25 userdata keys — Linux

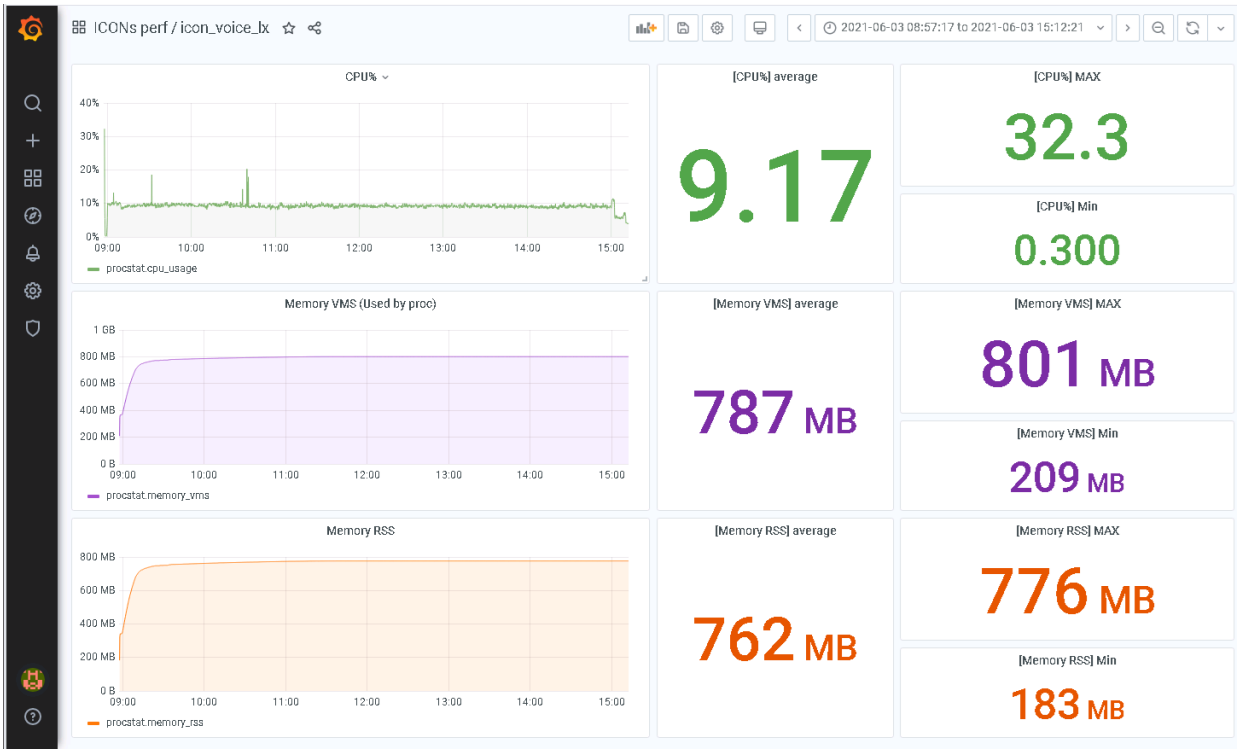

Figure 7: Voice baseline performance test — 10 userdata keys — Linux

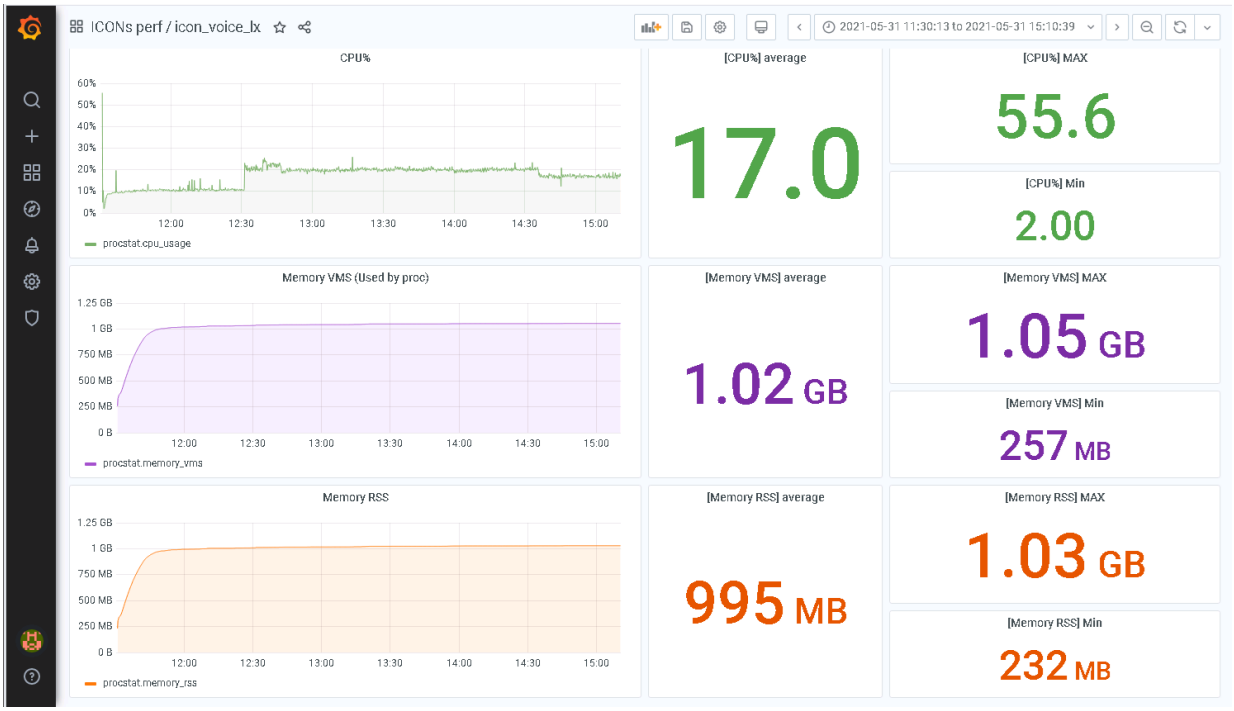

Figure 8: Voice baseline performance test with backlog — Linux

GENESYS

Chapter

3 Multimedia Tests

This chapter describes the Interaction Concentrator 8.1.5 Multimedia testing. The chapter contains the following sections:

Multimedia Test Environment

The multimedia test environment was organized as one tenant containing two Switches and two Interaction Servers. Both Interaction Servers were served by a single ICON and IDB.

The configuration contained:

- 40,000 Persons
- 40,000 Places
- 100 Simple Routing Strategies
- 100 Interaction Queues

A total of 8,000 agents were logged in.

Call configuration included attached data with 10 key-value pairs (KVPs) as described under Load specification for Multimedia.

Tests were conducted independently on Windows and Linux. A single ICON monitored all activity and stored data in a single IDB. ICON options to filter the data that is stored in IDB were turned off.

Each test started with an empty IDB and empty persistent queue file (empty backlog). IDB was not purged during the test.

Two database connection scenarios were used:

- Database connection is stable during the test (empty backlog)

- Database connection is lost during the test (non-empty backlog):
  - For the first hour of the test and under a normal workload, loss of database connection is simulated. The persistent queue file grows for the hour that the database connection is broken.
  - Then database connection is restored.

Load specification for Multimedia

Interaction (call) parameters for ICON load:

- 15.6 calls per second (cps)
- 2 Interaction Servers
- 8,000 Agent login sessions

Userdata parameters:

- 10 keys
- Length of key names: 32 bytes
- Length of key values: 124 bytes

The running load interval was 6 hours. The interaction load was produced by two Interaction Servers. Approximately 340,000 multimedia interactions in total were processed within the interval.

Multimedia interaction lifespan was set randomly from 60 to 600 seconds. The lifespan of login sessions ranged from 180 to 3,600 seconds (3 minutes to 1 hour).

The following table shows the distribution of interaction scenarios.

Interaction (chat) scenario	% of total calls
Simple Chat Short	90%
Simple Chat Long	2%
Chat with Transfer	2%
Chat with Conference	2%
Agent Reject Invite	2%
Invite expired	2%

Multimedia Performance Test Results

Results Summary

The following table summarizes the multimedia performance test results for the baseline and backlog tests.

Load specification	Application	Average CPU	Max CPU	Average RAM	Max RAM
Multimedia Baseline performance test • 8,000 login sessions • Call rate of 15.6 cps					
• 10 userdata keys	ICON (Windows)	10.6%	18.7%	713 MB	762 MB
	ICON (Linux)	11.4%	16.6%	727 MB	748 MB
Multimedia Baseline performance test with backlog • 8,000 login sessions • Call rate of 15.6 cps • For the first hour ICON DB Servers (primary and backup) are stopped to prevent ICON from storing data in the database					
• 10 userdata keys	ICON (Windows)	10.9%	24.4%	711 MB	760 MB
	ICON (Linux)	10.8%	21.8%	725 MB	748 MB

Grafana Plots for Multimedia Performance Tests

This section shows Multimedia performance test data visualized through Grafana for Windows and Linux, respectively.

- Multimedia Performance Tests on Windows
- Multimedia Performance Tests on Linux

Multimedia Performance Tests on Windows

The screenshots in this section show Grafana graphs and metrics for ICON 8.1.5 Multimedia baseline tests on Windows with 10 userdata keys, with and without a backlog.

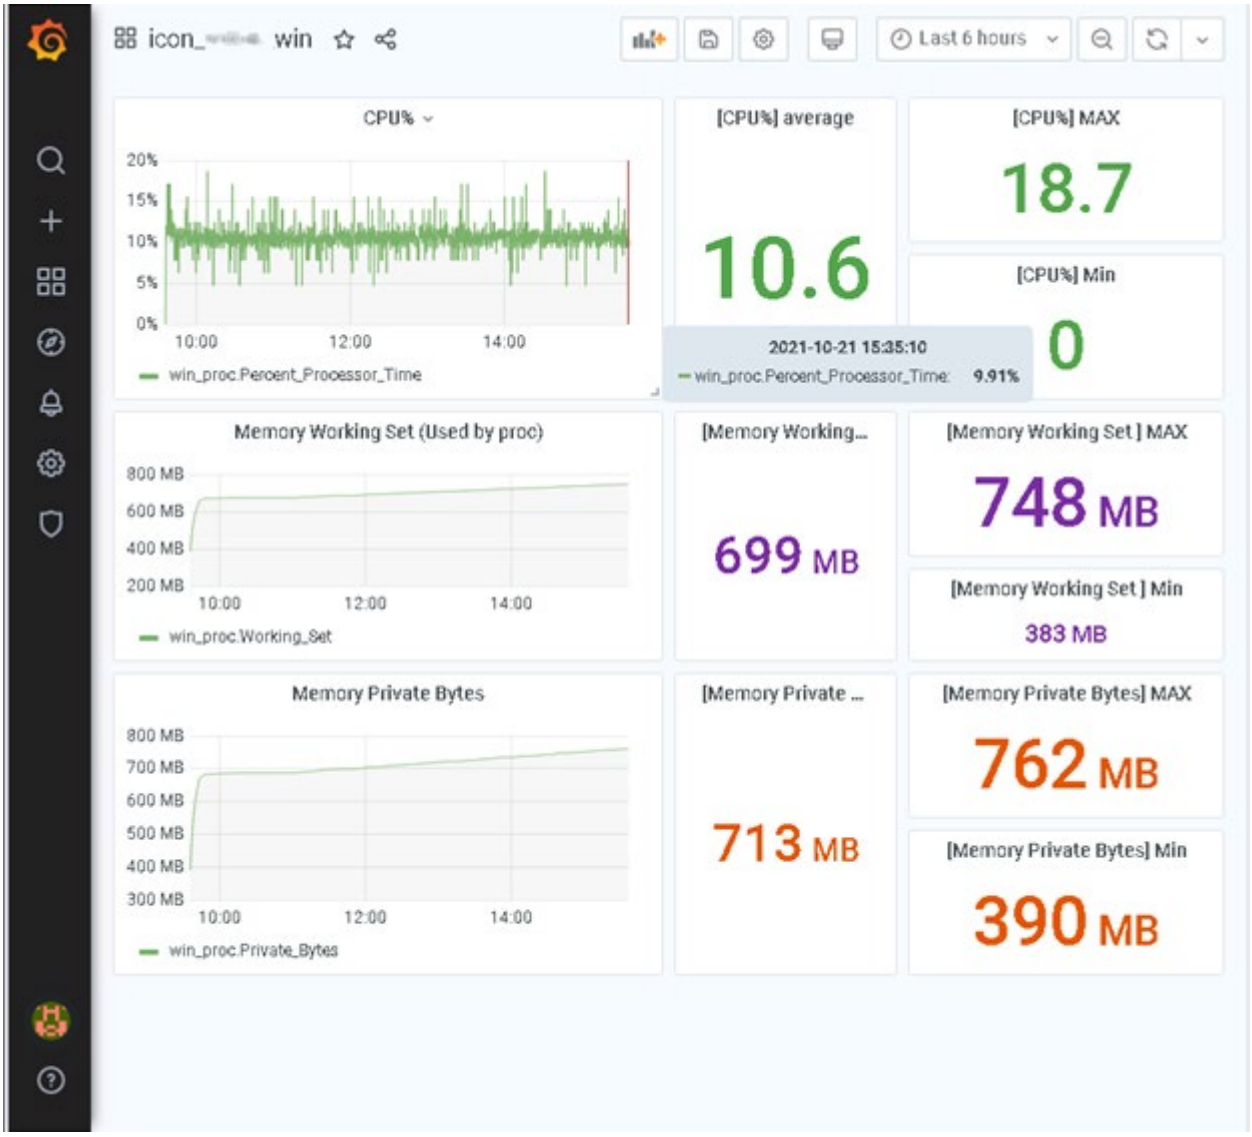

Figure 9: Multimedia baseline performance test — Windows

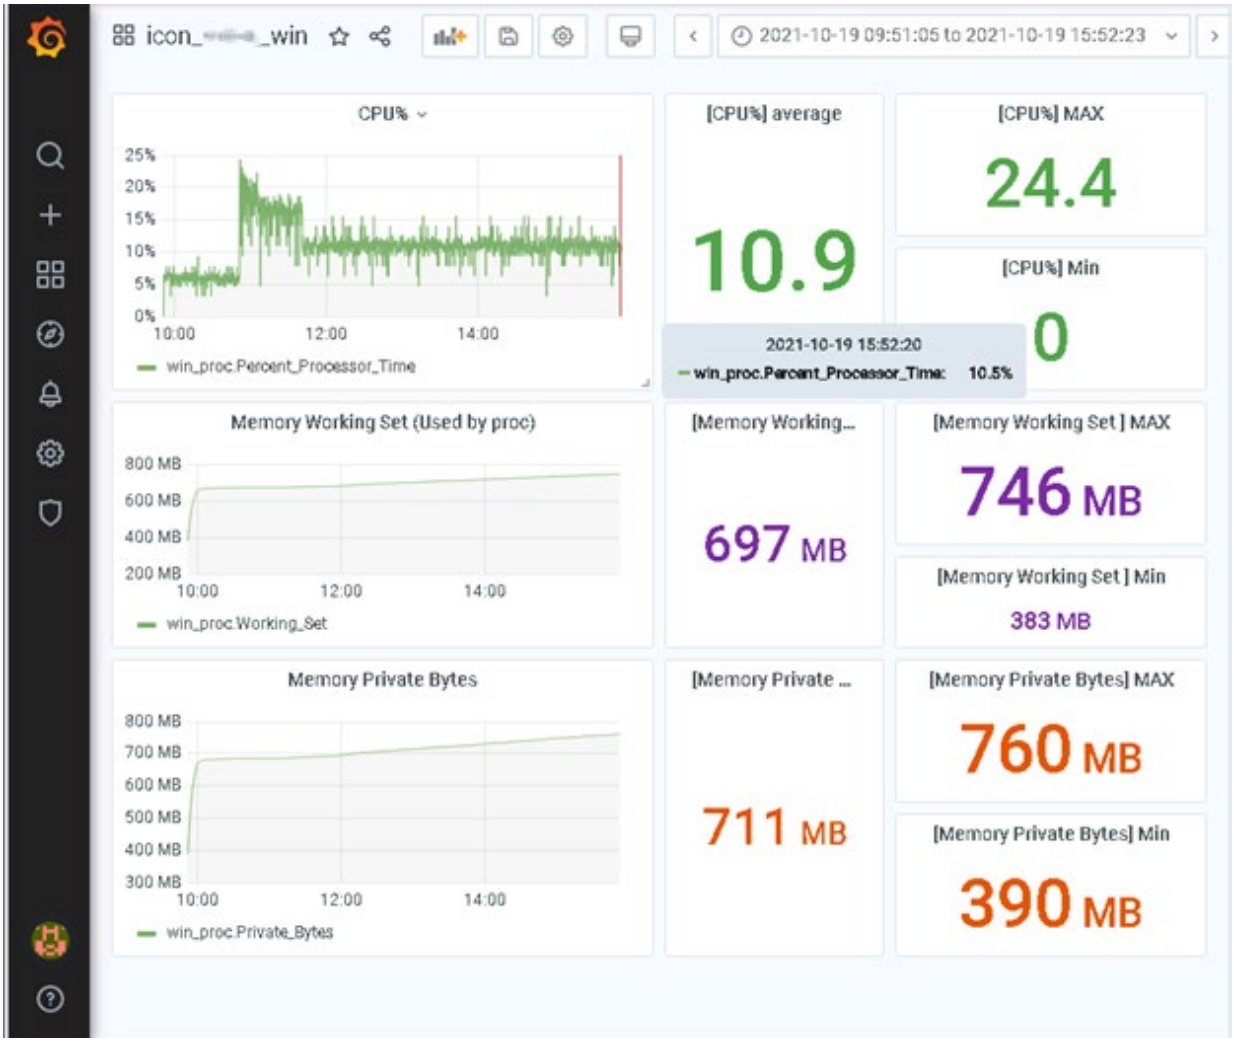

Figure 10: Multimedia baseline performance test with backlog — Windows

Multimedia Performance Tests on Linux

The screenshots in this section show Grafana graphs and metrics for ICON 8.1.5 Multimedia baseline tests on Linux with 10 userdata keys, with and without a backlog.

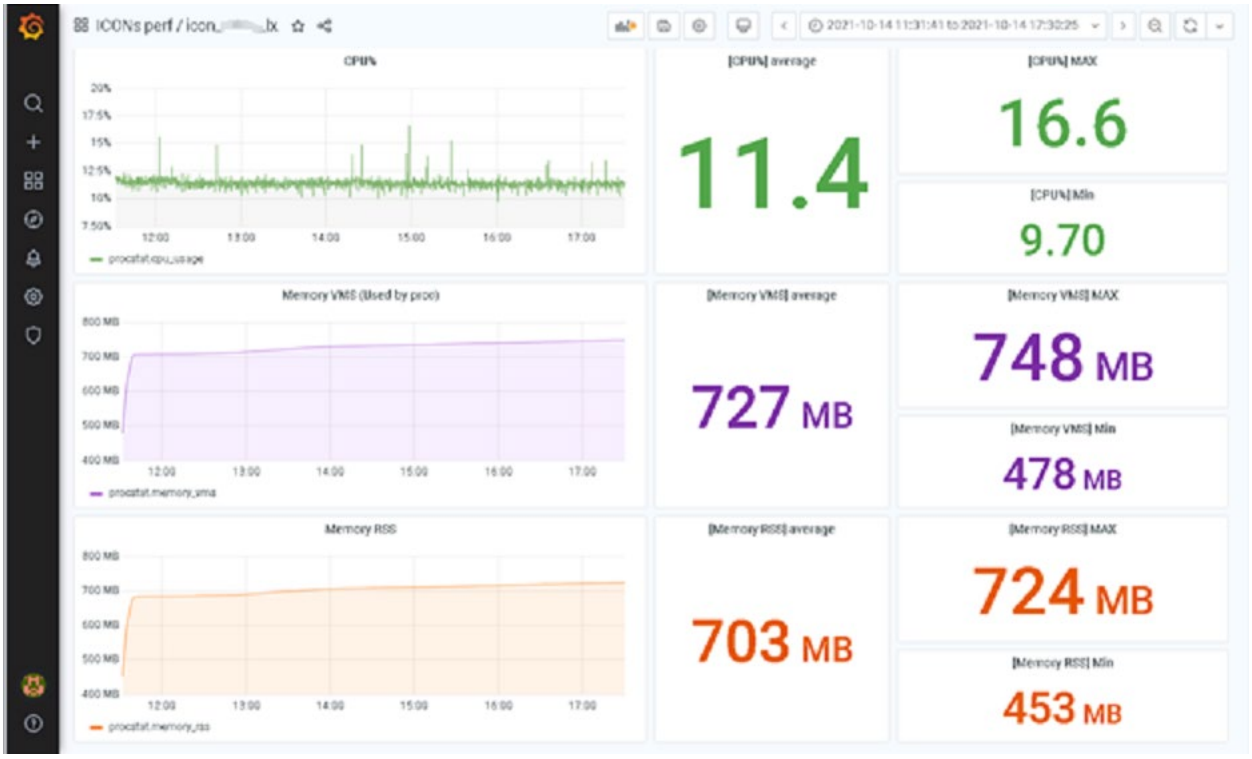

Figure 11: Multimedia baseline performance test — Linux

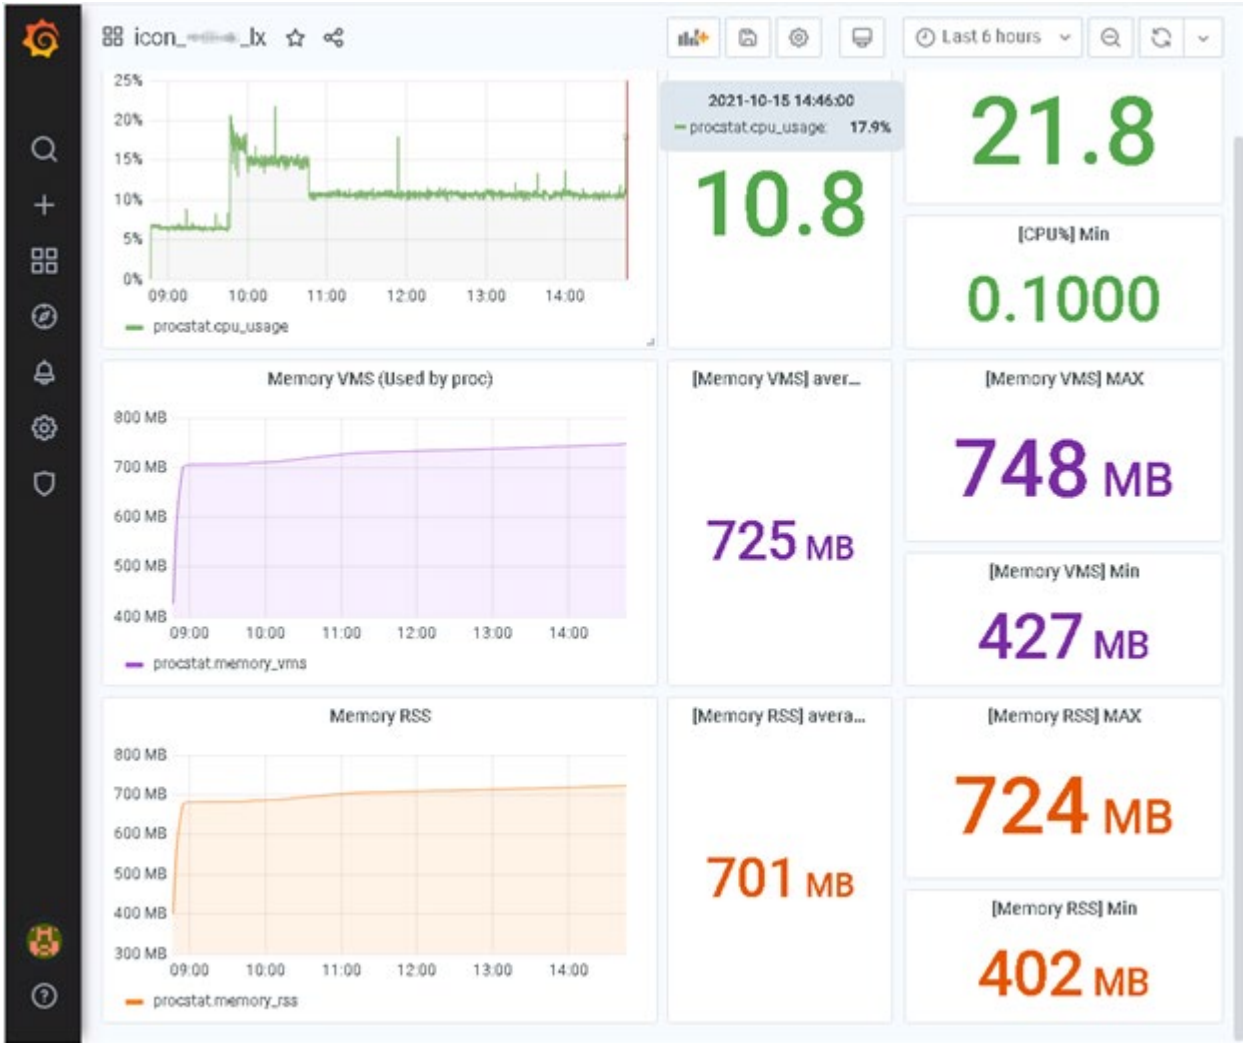

Figure 12: Multimedia baseline performance test with backlog — Linux

GENESYS

Chapter

4 ICON 8.1.5 Performance Conclusions

Based on the results of performance testing that was conducted in a large-scale environment, Genesys has the following observations and conclusions:

- ICON can recover successfully if a backlog builds up under load conditions.
- The memory management options are important in controlling ICON memory consumption, so that ICON can continue to operate under load conditions while a large number of interactions are backlogged.