
Artículo

[Mark Bolinsky](#) · 6 mar, 2020 Lectura de 7 min

InterSystems IRIS for Health 2020.1 HL7 Benchmark

Introduction

InterSystems has recently completed a performance and scalability benchmark of IRIS for Health 2020.1, focusing on HL7 version 2 interoperability. This article describes the observed throughput for various workloads, and also provides general configuration and sizing guidelines for systems where IRIS for Health is used as an interoperability engine for HL7v2 messaging.

The benchmark simulates workloads that have been designed to closely match live environments. The details of the simulation are described in the Workload Description and Methodology section. The tested workloads comprised HL7v2 Patient Administration (ADT) and Observation Result (ORU) payloads and included transformations and re-routing.

The 2020.1 version of IRIS for Health has demonstrated a sustained throughput of over 2.3 billion (total inbound and outbound) messages per day with a commodity server using the 2nd Generation Intel® Xeon® Scalable Processors and Intel® Optane™ SSD DC P4800X Series SSD storage. These results have more than doubled the scalability from the prior Ensemble 2017.1 HL7v2 throughput benchmarking.

Throughout these tests, IRIS for Health was configured to preserve first-in/first-out (FIFO) ordering, and to fully persist messages and queues for each inbound and outbound message. By persisting the queues and messages, IRIS for Health provides data protection in the event of a system crash, and full search and resend capabilities for historic messages.

Further, configuration guidelines are discussed in the sections below, which will assist you in choosing an appropriate configuration and deployment to adequately meet your workload 's performance and scalability requirements.

The results demonstrate that IRIS for Health is capable of satisfying extreme messaging throughput on commodity hardware, and in most cases allowing a single small server to provide HL7 interoperability for an entire organization.

Overview of Results

Three workloads were used to represent different aspects of HL7 interoperability activity:

- **T1 workload:** uses simple pass-through of HL7 messages, with one outbound message for each inbound message. The messages were passed directly from the Ensemble Business Service to the Ensemble Business Operation, without a routing engine. No routing rules were used and no transformations were executed. One HL7 message instance was created in the database per inbound message.
- **T2 workload:** uses a routing engine to modify an average of 4 segments of the inbound message and route it to a single outbound interface (1-to-1 with a transform). For each inbound message, one data transformation was executed and two HL7 message objects were created in the database.
- **T4 workload:** uses a routing engine to route separately modified messages to each of four outbound interfaces. On average, 4 segments of the inbound message were modified in each transformation (1 inbound to 4 outbound with 4 transforms). For each inbound message four data transformations were executed, four messages were sent outbound, and five HL7 message objects were created in the database.

The three workloads were run on a physical 48-core system with two Intel® Scalable Gold 6252 processors with

two 750GB Intel® Optane™ SSD DC P4800X SSD drives running Red Hat Enterprise Linux 8. The data is presented as the number of messages per second (and per hour) inbound, the number per second (and per hour) outbound, as well as the total messages (inbound plus outbound) in a 10-hour day. Additionally, CPU utilization is presented as a measure of available system resources at a given level of throughput.

Scalability Results

Table-1: Summary of throughput of the four workloads on this tested hardware configuration:

Test	# of Inbound Business Services	Inbound		Outbound		Total In+Out messages/day (10 hours)	Total In+Out messages/day (24 hours)	CPU Utilization
		Messages Per Second	Messages Per Hour	Messages Per Second	Messages Per Hour			
T1	64	15,997	57,589,200	15,997	57,589,200	1,151,784,000	2,764,281,600	78%
T2	48	8,891	32,007,600	17,782	64,015,200	960,228,000	2,304,547,200	79%
T4	128	5,399	19,436,400	21,596	77,745,600	971,820,000	2,332,368,000	73%
* T1 + T2 + T4	16 + 16 + 32	7,019	25,268,400	14,038	50,535,551	758,039,511	1,819,294,826	70%

* Combined workload with 25% of T1 / 25% of T2 / 50% T4 workload mix

Workload Description and Methodology

The tested workloads included HL7v2 Patient Administration (ADT) and Observation Result (ORU) messages, which had an average size of 1.2KB and an average of 14 segments. Roughly 4 segments were modified by the transformations (for T2 and T4 workloads). The tests represent 48 to 128 inbound and 48 to 128 outbound interfaces receiving and sending messages over TCP/IP.

In the T1 workload, four separate namespaces each with 16 interfaces were used, and the T2 workload used three namespaces each with 16 interfaces, the T4 workload used four namespaces each with 32 interfaces, and the final “mixed workload” used three namespaces with 16 for T1 workload, 16 for T2 workload, and 32 for T4 workload.

The scalability was measured by gradually increasing traffic on each interface to find the highest throughput with acceptable performance criteria. For the performance to be acceptable the messages must be processed at a sustained rate, with no queuing, no measurable delays in delivery of messages and the average CPU usage must remain below 80%.

Previous testing has demonstrated that the type of HL7 message used is not significant to the performance or scalability of Ensemble; the significant factors are the number of inbound messages, the size of inbound and outbound messages, the number of new messages created in the routing engine, and the number of segments modified.

Additionally, previous testing has shown that processing individual fields of an HL7 message in a data transformation is not usually significant to performance. The transformations in these tests used fairly straightforward assignments to create new messages. Note that complex processing (such as use of extensive SQL queries in a data transformation) may cause results to vary.

Previous testing has also verified that rules processing is not usually significant. The routing rule sets used in these tests averaged 32 rules, with all rules being simple. Note that extremely large or complex rule sets may cause results to vary.

Hardware

Server Configuration

The tests utilized a server with 2nd Generation Intel® Scalable Gold 6252 “Cascade Lake” processors providing 48

cores @ 2.1GHz on a 2-socket system, 24 cores per socket with 192 GB DDR4-2933 DRAM, and 10Gb Ethernet network interface. Red Hat Enterprise Linux Server 8 operating system was used for this test with InterSystems IRIS for Health 2020.1

Disk Configuration

Messages passing through IRIS for Health are fully persisted to disk. In the case of this test two Intel 750GB Intel® Optane™ SSD DC P4800X SSD drives internal to the system were used splitting the databases on one drive and the journals on another. In addition to ensure real-world comparison synchronous commit is enabled on the journals to force data durability. For the T4 workload as described previously in this document, each inbound HL7 message generates roughly 50KB of data, which can be broken down as described in Table 2. Transaction journals are typically kept on line for less time than message data or logs and this should be taken into account when calculating the total disk space required.

Table 2: Disk Requirement per inbound HL7 T4 Message

Contributor	Data Requirement
Segment Data	4.5 KB
HL7 Message Object	2 KB
Message Header	1.0 KB
Routing Rule Log	0.5 KB
Transaction Journals	42 KB
Total	50 KB

Recall that the T4 workload used a routing engine to route separate modified messages to each of four outbound interfaces. On average, 4 segments of the inbound message were modified in each transformation (1 inbound to 4 outbound with 4 transforms). For each inbound message four data transformations were executed, four messages were sent outbound, and five HL7 message objects were created in the database.

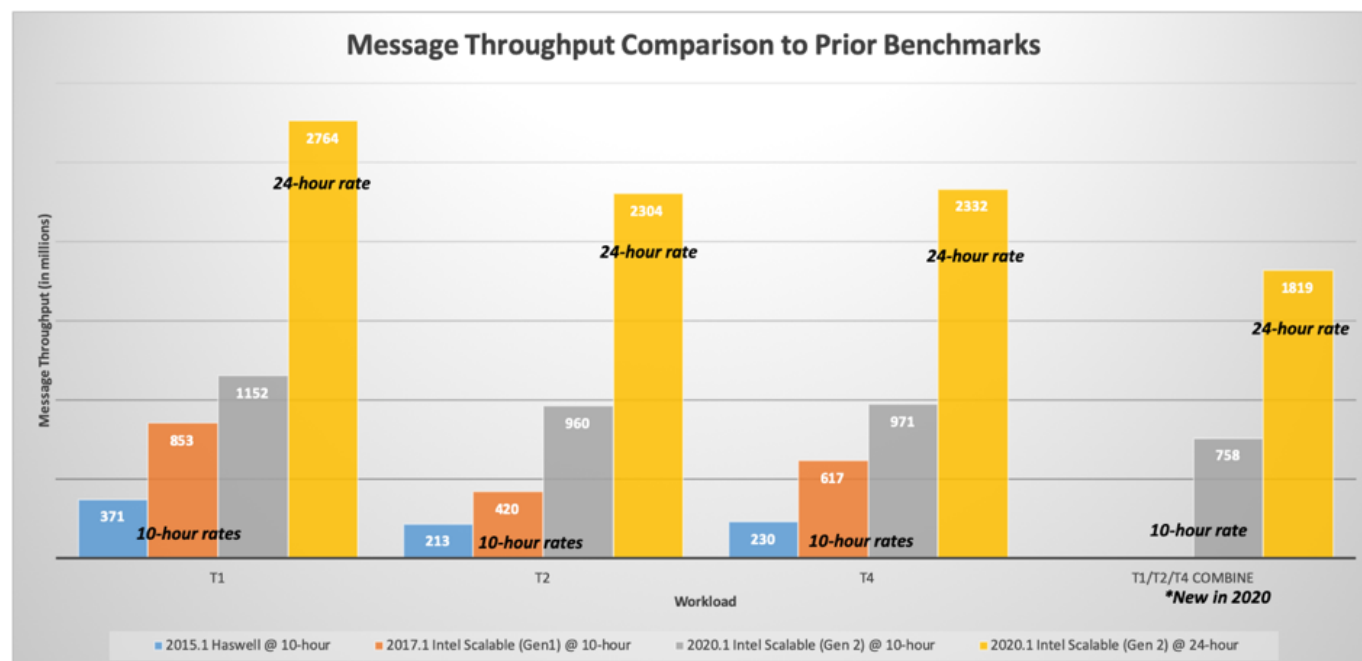
When configuring systems for production utilization, net requirements should be calculated by considering the daily inbound volumes as well as the purging schedule for HL7 messages and the retention policy for journal files. Additionally, appropriate journal file space should be configured on the system so as to prevent the journal disk volumes from filling up. The journal files should reside on physically separate disk from the database files, for both performance as well as reliability considerations.

Conclusion

InterSystems IRIS for Health HL7v2 message throughput demonstrated in these tests illustrates the massive throughput capabilities with a modest 2-socket commodity server configuration to support the most demanding message workloads of any organization. Additionally, InterSystems is committed to constantly improving on the performance and scalability from version to version along with taking advantage of the latest server and cloud technologies.

The following graph provides an overview and comparison of the increase in throughput from the previous Ensemble 2015.1 and Ensemble 2017.1 benchmarks with the Intel® E5-2600 v3 (Haswell) processors and Ensemble 2017.1 benchmark with 1st Generation Intel® Scalable Platinum Series (Skylake) processors respectively to latest results with the 2nd Generation Intel® Scalable Gold Series (Cascade Lake) processors running IRIS for Health 2020.1.

Graph-1: Message throughput (in millions) per 10-hour day on a single server



InterSystems IRIS for Health continues to raise the bar on interoperability throughput from version to version along with offering flexibility in connectivity capabilities. As the above graph shows the message throughput has increased significantly and, in case of T2 workloads, doubled from 2017, and comparing to 2015 more than tripled throughput in the same 10-hour window and sustained over 2.3 billion total 24-hour message rates.

Another key indicator of the advancements of IRIS for Health is the throughput improvement in the more complex T2 and T4 workloads which incorporates transformations and routing rules as opposed to pure pass-through operation of the T1 workload.

[InterSystems is available to discuss solutions for your organizations interoperability needs.](#)

[#HL7](#) [#Interoperabilidad](#) [#Arquitecturas y Soluciones de Negocio con InterSystems](#) [#Rendimiento](#) [#InterSystems IRIS](#) [#InterSystems IRIS for Health](#)

URL de fuente: <https://es.community.intersystems.com/node/474686>