

## INTERSYSTEMS ENSEMBLE® HL7v2 MESSAGE THROUGHPUT

*Ensemble (v 2010.2, build 503) HL7v2 Performance & Scalability; December, 2010*

# INTERSYSTEMS ENSEMBLE® HL7v2 MESSAGE THROUGHPUT

## Executive Summary

InterSystems Ensemble® is a rapid integration and development platform with built-in capabilities for high-speed processing of HL7 messages. It has been ranked as the #1 or #2 Interface Engine in healthcare every year since 2006 by KLAS, a prominent healthcare IT analyst organization based in the United States.

InterSystems recently completed a performance and scalability benchmark of Ensemble version 2010.2 (Maintenance Release 1), focusing on Health Level 7 version 2 (HL7v2) messaging. This document describes the observed characteristics, and also provides general configuration and sizing guidelines for systems where Ensemble is used as an interface 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 were comprised of HL7v2 Patient Administration (ADT) and Observation Result (ORU) payloads, and included transformations and re-routing.

Inbound			Outbound			Total (in + out)
Per Second	Per Hour	Per Day (10 hours)	Per Second	Per Hour	Per Day (10 hours)	Per Day (10 hours)
326	1,173,600	11,736,000	1,304	4,694,400	46,944,000	<b>58,680,000</b>

TABLE 1: DAILY SUSTAINABLE HL7 MESSAGE RATES WITH ENSEMBLE 2010.2

As Table 1 indicates, an excess of 58 million combined inbound and outbound messages were demonstrated as sustainable per day (over 10 hours) on the tested 12-core system. This workload is described later in this document as the T4 workload, and included a routing engine that sends separate modified messages to each of four destinations.

Throughout these tests, Ensemble was configured to preserve first-in/first-out order, and to fully persist messages and queues for each inbound and outbound message. By persisting the queues and messages, Ensemble provides data protection in the event of a system crash, and full search and resend capabilities for all messages.

Further, configuration guidelines are discussed in this document's appendix, 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 Ensemble is capable of supporting large-scale messaging needs on fairly small (commodity) hardware, allowing a single small server to potentially provide HL7 messaging support for the entire organization.

<sup>1</sup> www.KLASresearch.com/top\_20.

## Overview of Results

Three workloads were used to represent different aspects of Ensemble activity:

- The T1 workload used 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.
- The T2 workload used 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.
- 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-to-4 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 at the 6- and 12-core levels. 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. 16 inbound and 16 outbound interfaces were used in all cases.

### 6-CORE SCALABILITY

Table 2 summarizes the best scenarios in the 6-core configuration:

Test	Inbound		Outbound		CPU Utilization	Total (in + out) messages/day (10 hours)
	Per Second	Per Hour	Per Second	Per Hour		
T1	616	2,217,600	616	2,217,600	34%	<b>44,352,000</b>
T2	332	1,195,200	332	1,195,200	34%	<b>23,904,000</b>
T4	144	518,400	576	2,073,600	35%	<b>25,920,000</b>

TABLE 2: 6-CORE ENSEMBLE 2010.2 HL7V2 SCALABILITY

### 12-CORE SCALABILITY

Table 3 summarizes the best scenarios for each of the three workloads in the 12-core configuration:

Test	Inbound		Outbound		CPU Utilization	Total (in + out) messages/day (10 hours)
	Per Second	Per Hour	Per Second	Per Hour		
T1	925	3,330,000	925	3,330,000	25%	<b>66,600,000</b>
T2	678	2,440,800	678	2,440,800	37%	<b>48,816,000</b>
T4	326	1,173,600	1,304	4,694,400	38%	<b>58,680,000</b>

TABLE 3: 12-CORE ENSEMBLE 2010.2 HL7V2 SCALABILITY

Finally, as Figure 1 illustrates, Ensemble can scale almost linearly for each type of test when the number of cores is doubled from 6 to 12.




FIGURE 1: ENSEMBLE 2010.2 SCALABILITY PROFILES ACROSS 6- AND 12-CORE CONFIGURATIONS

## 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 16 inbound and 16 outbound interfaces receiving and sending messages over TCP/IP. Scalability was measured by gradually increasing traffic on each interface. Tests varying the number of interfaces made little difference to the overall throughput or performance, which is why 16 inbound and outbound interfaces were used for these tests.

An acceptable test demonstrated sustained steady state with minimal queuing. Queuing is indicative of a poor quality-of-service for messages, and can result in a delay of processing and delivery of a message to the destination system(s). The test methodology rejected any tests with queuing that would result in a single message being delayed for over 1 second. Additionally, tests that resulted in CPU utilization greater than 75% were generally rejected. The reported results in this document also include a benchmarking discount (20%) to account for variances between lab testing and live production systems.

Previous testing 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 showed 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 resource-intensive SQL queries in a data transformation) may cause results to vary.

Previous testing 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. 'Do All Rules' was set to false, and all rules in the rule sets were exercised equally often. Note that extremely large or complex rule sets may cause results to vary.

## Hardware Used

The tests utilized a server with dual hexa-core Intel Xeon X5670 "Westmere" processors (12 cores @ 2.93GHz, 2 chips, 6 cores per chip, 2 threads per core), with 32 GB RAM. The OS used was Red Hat Enterprise Linux Server release 5.5 (Tikanga).

HyperThreading (HT) technology was enabled for these tests. There were a series of tests conducted to verify the impact of HT on Ensemble performance. The conclusion was that HT on the Intel Xeon 5670 (Westmere) processor provides a net positive impact on performance by decreasing the amount of time taken to complete a workload. Please note that HT on other architectures and chips might have a negative (or neutral) impact on performance, and care should be taken to evaluate the performance on those platforms prior to enabling HT. Additionally, a relatively modest disk configuration was used for this test where the Database Files, Journal Files, and Write Image Journal (WIJ) were all configured on one internal disk for simplicity (not recommended for production systems).

## Appendix: Configuration Guidelines

In addition to standard InterSystems platform-specific configuration and tuning guidelines for InterSystems Cache® and Ensemble, the following guidelines apply.

### CPU CONFIGURATION

The Standard Performance Evaluation Corporation (SPEC) provides standardized benchmarks that can provide a baseline for configuration across platforms. One of the components of the SPEC benchmarks is CINT2006, which tests the integer processing performance of a given CPU. SPEC defines a base runtime for the benchmark programs, which is known as the reference time. Timed tests are then run on various test systems, and those numbers are compared to the reference time and a ratio is computed. That ratio is then referred to as the SPECint2006 score (also known as CINT2006 score) for that test. The SPECint score typically is used to compare cross-platform configurations.

The system we tested was configured with two hexa-core Intel Xeon X5670 processors (2.93 GHz). This processor has a CINT2006 baseline score of 36.6, which is the score for one processing core. Therefore, the 12-core configuration would have a CINT2006 baseline score of 439.2 (36.6 x 12). These numbers can be used as a guideline to compare other platforms and determine the amount of computing power they require to sustain the levels achieved on the Intel Xeon X5670 platform.

For example, the Intel Xeon X5550 (2.66 GHz) processor has a CINT2006 baseline score of 29.7, which indicates that its throughput is roughly 19% lower (29.7/36.6) than the tested Intel Xeon X5670 system. Therefore, if you configure a comparable system with the Intel Xeon X5550 processor instead of the Intel Xeon X5670 processor, you would expect Ensemble to process roughly 19% fewer messages.

### DISK CONFIGURATION

As previously mentioned, messages passing through Ensemble are fully persisted to disk. 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 on the next page.

Contributor	Data Requirement
Segment Data	4.5 KB
HL7 Message Object	2 KB
Message Header	1.0 KB
Routing Rule Log	0.5 KB
Journaling Information (in Caché Journal File)	42 KB
Total	50 KB

TABLE 4: DISK REQUIREMENT PER INBOUND HL7 T4 MESSAGE

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-to-4 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. Additionally, appropriate journal file space should be configured on the system to prevent the journals from filling up. The journal files should reside on physically separate disk from the database files, for performance and reliability considerations.

---

InterSystems Corporation

World Headquarters

One Memorial Drive

Cambridge, MA 02142-1356

Tel: +1.617.621.0600

Fax: +1.617.494.1631

[InterSystems.com](http://InterSystems.com)

**INTERSYSTEMS**