*An **InterSystems** Technical White Paper*

One Memorial Drive, Cambridge, MA 02142, USA
Tel: +1.617.621.0600
www.intersystems.com

InterSystems Ensemble® and InterSystems HealthShare ® HL7v2 Message Throughput

Ensemble (v 2015.1, build xxx) HL7v2 Performance & Scalability; March 2015

CONTENTS

By: Mark Bolinsky, Technology Architect Manager, and
David Loveluck, Director of Product Management

EXECUTIVE SUMMARY

InterSystems HealthShare® and InterSystems Ensemble® both provide a rapid integration and development platform with built-in capabilities for the high-speed processing of HL7 messages. For the purposes of HL7 v2 message routing the two products are equivalent in performance. For brevity, this document will just say Ensemble in many places but it should be taken to apply equally to both products.

We have recently completed a performance and scalability benchmark of Ensemble version 2015.1, focusing on HL7 version 2 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 comprised HL7v2 Patient Administration (ADT) and Observation Result (ORU) payloads and included transformations and re-routing.

Ensemble was demonstrated to sustain a throughput of over 200 million (total inbound plus outbound) messages per 10-hour day on the tested 28-core system with adequate headroom for further scaling.

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 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 Ensemble is capable of satisfying extreme messaging needs on commodity hardware, potentially allowing a single small server to provide HL7 messaging support for an entire organization.

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 separately 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 on a 28-core Intel E5-2697 v3 based system running Red Hat Linux. 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. In all cases, 32 inbound and 32 outbound interfaces were used.

SCALABILITY

Table 1 summarizes the best scenarios for each of the three workloads at a 28-core configuration:

	Inbound		***Outbound***			
Test	*Per Second*	*Per Hour*	*Per Second*	*Per Hour*	*Total (in + out) messages/day (10 hours)*	*CPU Utilization*
T1	5,151	18,543,600	5,151	18,543,600	370,872,000	32%
T2	2,953	10,630,800	2,953	10,630,800	212,616,000	43%
T4	1,280	4,608,000	5,120	18,432,000	230,400,000	43%

Table 1: 28-core Ensemble 2015.1 HL7v2 scalability

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 32 inbound and 32 outbound interfaces receiving and sending messages over TCP/IP. 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 significant queuing, no measurable delays in delivery of messages and the average CPU usage must remain below 75%. The reported numbers have been reduced by a 20% benchmark discount in line with common practice to allow for statistical variations and other factors that might arise in a synthetic test environment.

Earlier tests with much larger numbers of interfaces made little difference to the overall throughput or performance, which is why 32 inbound and outbound interfaces were used for these tests.

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 expensive 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. Each rules in the rule sets were equally often. Note that extremely large or complex rule sets may cause results to vary.

HARDWARE USED

The tests utilized a server with dual 14-core Intel Xeon E5-2697 v3 "Haswell" processors providing 28 cores @ 2.6GHz on a 2-socket system, 14 cores per chip, with 64 GB RAM, and 800GB internal SSD disk. Red Hat Enterprise Linux Server 7 operating system was used for this test.

In addition to standard InterSystems platform-specific configuration and tuning guidelines for Caché and Ensemble, the following guidelines apply:

CPU CONFIGURATION

The Standard Performance Evaluation Corporation (SPEC) provides standardized benchmarks that can be run across various platforms and configurations in order to provide a baseline for configuration across platforms. One of the components that SPEC benchmarks test is the CINT2006, which is the integer component of the SPEC benchmarks. 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 can typically be used to compare cross-platform configurations.

The system under test was configured with two 14-core Intel Xeon E5-2697 v3 processors (2.6 GHz). This processor has a CINT2006 base of 63.4, which is the standardized performance score for one single processing core. The 28-core configuration has a CINT2006 rate baseline score of 1190, which is the standardized scalability score for all 28-cores. These numbers can be used as a guideline to compare other platforms, so as to determine the amount of computing power required on those platforms to sustain the levels achieved on the Intel Xeon E5-2697 v3 platform.

For example, the Intel Xeon E5-2650 v3 (2.30 GHz) processor has a CINT2006 baseline of 54.1, which indicates that its throughput is approximately 15% lower (54.1/63.4) than the tested Intel Xeon E5-2697 v3 system. Therefore, if you configure a system with the Intel Xeon E5-2650 v3 processor @2.30 Ghz instead of the Intel Xeon E5-2697 v3 processor @ 2.60 Ghz, one would expect that Ensemble would process approximately 15% fewer messages for comparable configurations and a higher reported percentage CPU utilization due to fewer processors cores per chip.

DISK CONFIGURATION

Messages passing through Ensemble are fully persisted to disk. In the case of this test a 800GB SSD disk internal to the system was used for both database and journal files. 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 4. 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.

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

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 and the retention policy for journal files. Additionally, appropriate journal file space should be configured on the system so as to prevent the journals from filling up. The journal files should reside on physically separate disk from the database files, for both performance as well as reliability considerations.

ABOUT HEALTHSHARE

InterSystems HEALTHSHARE®

Used worldwide, InterSystems HealthShare® is the premier informatics platform and solutions set for connected healthcare, enabling organizations to capture and share all patient data, with the interoperability and real-time analytics to drive informed action across a hospital network, community, region, or an entire nation.

HealthShare Health Connect provides the integration functionality of Ensemble plus support for additional healthcare protocols and document formats.

For more information, visit InterSystems.com/HealthShare.

ABOUT ENSEMBLE

InterSystems Ensemble® is a seamless platform for integration and the development of new connected applications. It is used by enterprises that want a rapid integration platform, and by application providers who want to create breakthrough "*connectable*" applications that can share data with other applications. Ensemble also leverages previous software investments through composite applications, and establishes an enterprise service bus (ESB) or SOA infrastructure.

For enterprises with challenging integration tasks, Ensemble is the easiest integration platform to use because it's not a stitched-together suite of separate parts. We created it as a single, architecturally consistent technology stack (integration server, data server, application server, and portal development software). So Ensemble projects are typically completed in half the time required with previous generations of integration products.

For application providers, the Ensemble platform enables the rapid development of new connectable applications that have embedded integration capability. In addition, Ensemble makes it easy for developers to enhance existing products with valuable features such as adaptable workflow, browser-based user interfaces, dashboards and rules-based business processes – without rewriting code.

For more information, visit InterSystems.com/Ensemble.

ABOUT INTERSYSTEMS

InterSystems develops advanced software technologies that enable breakthroughs. With a passion for excellence and a focus on client success, InterSystems provides data management, strategic interoperability, and analytics platforms used in healthcare, financial services, government, and dozens of other industries. InterSystems also offers unified healthcare applications, based on its core technologies, that deliver on the promise of connected healthcare. Founded in 1978, InterSystems is a privately held company headquartered in Cambridge, Massachusetts (USA), with offices worldwide. Its products are used daily by millions of people in more than 100 countries.

For more information, visit InterSystems.com.