

**Albert-Battaglin Consulting Group TAGITT-CATIA
V4/ENOVIA DMU Evaluation**

Table of Contents

Notices	2
Executive Summary	3
Summary of Top Three Winners in Each Test:	5
Individual Application Test Throughput Results:.....	7
Why TAGITT?.....	9
The Need for Application-Level Testing	9
Test Description	10
TAGITT-CATIA/ENOVIA DMU	12
Test Weighting	14
Machines Evaluated	21
TAGITT-CATIA/ENOVIA DMU Results	23
Conclusions.....	31
Methodology	33
Appendix.....	34

**Albert-Battaglin Consulting Group TAGITT-CATIA
V4/ENOVIA DMU Evaluation**

Table of Figures

Figure 1 – Some of the CATIA models used for TAGITT/CATIA	12
Figure 2 – ENOVIA DMU Navigator Test Models.....	13
Figure 3 – V4 Overall Throughput Weighting	15
Figure 4 – V4 Application Weighting Factors.....	16
Figure 5 – V4 Graphics Weighting Factors	17
Figure 6 – V4 Responsiveness Weighting Factors	18
Figure 7 – ENOVIA DMU Navigator Weighting Factors.....	19
Figure 8 – ENOVIA DMU Graphics Navigator Weighting Factors	20

Albert-Battaglin Consulting Group TAGITT-CATIA V4/ENOVIA DMU Evaluation

Table of Result Charts

Chart 1 – Overall V4 Throughput.....	23
Chart 2 – V4 Graphics Throughput	24
Chart 3 – V4 Application Throughput.....	25
Chart 4 – V4 System Responsiveness.....	26
Chart 5 – V4 Dynamic Graphic Throughput	27
Chart 6 – ENOVIA DMU Navigator Overall Throughput	28
Chart 7 – ENOVIA DMU Navigator Graphics Throughput.....	29
Chart 8 – ENOVIA DMU Navigator Interference Check Throughput.....	30
Chart 9 – Solid Model Creation and Modification Throughput	34
Chart 10 – Finite Element Analysis (ANSOLID) Throughput.....	35
Chart 11 – NC STL Generation Throughput	36
Chart 12 – Detail View Generation Throughput.....	37
Chart 13 – Solid and Surface Analysis Throughput	38
Chart 14 – Model Save Throughput.....	39
Chart 15 – Walk Through Throughput	40
Chart 16 –Sheet Metal Part Development Throughput.....	41
Chart 17 – Fitting Simulation Throughput	42
Chart 18 – Kinematics Simulation Throughput	43
Chart 19 – Studio Throughput	44
Chart 20 – Image Viewer Throughput	45
Chart 21 – ENOVIA DMU Navigator CGR Creation Throughput	46
Chart 22 – ENOVIA DMU Navigator Load Throughput.....	47

TAGITT-CATIA V4/ENOVIA DMU Evaluation



albert-battaglin consulting group
one ironwood drive
soquel, ca 95073
+1 831-464-0600

October 2005

Notices

This document was researched and written by The Albert-Battaglin Consulting Group under a contract with IBM. The analysis, views, conclusion and recommendations in this document are those of The Albert-Battaglin Consulting Group and not necessarily those of IBM. This document is provided “AS IS” and may contain errors and inaccuracies. THE ALBERT-BATTAGLIN CONSULTING GROUP AND IBM DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. It is your responsibility to evaluate the information contained in this document and ensure that your use is in accordance with your company’s policies and practices including, in the case of IBM, the IBM Business Conduct Guidelines.

Executive Summary

The Albert-Battaglin Consulting Group evaluated eight engineering workstations from Hewlett-Packard, IBM and Sun using TAGITT (The Albert Group Interactive Throughput Test). TAGITT-CATIA V4/ENOVIA DMU measures workstation performance from a user perspective by simulating interactive work sessions using CATIA 4.2.4 refresh 02 and ENOVIA DMU R15 SP2 versions of this CAD/CAM application software package. Key results were as follows:

- All configurations of the IBM's new IntelliStation® POWER™ 285 workstation outperformed IBM IntelliStation POWER 275 configurations as well as the HP c8000 (2w) FireGL X3 and Sun Blade 2500 (1w) XVR-1200 by considerable margins.
- The IBM POWER 285 (2w) GXT4500P and IBM POWER 285 (2w) GXT6500P machines were nearly tied as the overall performance leaders in the TAGITT-CATIA V4 test. The IBM POWER 285 (1w) GXT6500P and IBM POWER 285 (1w) GXT4500P were 5% and 7% slower than the fastest machine respectively.
- The IBM POWER 285 (2w) GXT6500P was the clear winner in the Overall DMU Throughput measurement being 15% faster than the second place IBM POWER 285 (2w) GXT4500P due to the superior OpenGL graphics performance on the GXT6500P. The IBM POWER 285 (1w) GXT6500P and IBM POWER 285 (1w) GXT4500P machines were 21% and 35% slower than the fastest machine respectively.
- The HP c8000 (2w) FireGL X3 and IBM POWER 285 (2w) GXT4500P tied as the fastest V4 graphics workstations when graphics computation results were included. For pure dynamic graphic manipulations, the HP c8000 (2w) FireGL X3 was 29% faster than the IBM POWER 285 (2w) GXT4500P.
- The IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (1w) GXT6500P machines were the fastest for overall DMU graphic operations. These machines outperformed the second place IBM POWER 285 (1w) GXT4500P and IBM POWER 285 (2w) GXT4500P machines by more than 30%. The HP c8000 (2w) FireGL X3 was 58% slower than the leaders in this test.
- Machine performance was not the same across all tested applications. Although the 1.9 GHz IBM POWER 285 machines were the fastest in nearly all tests, the HP c8000 (2w) FireGL X3 also won in a few individual tests (Overall Dynamic Graphics Performance, and Walk Through). In addition, the relative performance of the machines varied significantly on the different tests.

- No advantage was apparent with HP's dual processor configuration. IBM's dual processor configurations performed significantly better than the uniprocessor configurations in several tests. Running V4, the dual processor configurations outperformed uniprocessor configurations in the Studio and Image Viewer tests by about 60% and 95% respectively. Running DMU navigator, the dual-processor configurations were 77% faster than their single processor cousins.
- The Sun Blade 2500 (1w) XVR-1200 performance was markedly worse than all other systems. In V4 overall throughput, the Sun Blade 2500 (1w) XVR-1200 was 139% slower than the IBM POWER 285 (1w) GXT6500P, IBM's fastest uniprocessor configuration. In ENOVIA DMU overall throughput, the Sun Blade 2500 (1w) XVR-1200 was 249% slower than the IBM POWER 285 (2w) GXT6500P due to extremely bad performance in the interference computation test.

Summary of Top Three Winners in Each Test:

V4 Throughput Summary Results			
Task/Test	First	Second	Third
Overall V4 Throughput	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (2w) GXT6500P	IBM POWER 285 (1w) GXT6500P IBM POWER 285 (1w) GXT4500P	IBM POWER 275 (1w) GXT6500P
Application Throughput	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (2w) GXT6500P IBM POWER 285 (1w) GXT4500P IBM POWER 285 (1w) GXT6500P	IBM POWER 275 (1w) GXT6500P IBM POWER 275 (1w) GXT4500P	HP c8000 (2w) FireGL X3
Graphics Throughput	HP c8000 (2w) FireGL X3 IBM POWER 285 (2w) GXT4500P	IBM POWER 285 (2w) GXT6500P IBM POWER 285 (1w) GXT6500P	IBM POWER 285 (1w) GXT4500P
Responsiveness Throughput	IBM POWER 285 (2w) GXT6500P IBM POWER 285 (2w) GXT4500P IBM POWER 285 (1w) GXT6500P IBM POWER 285 (1w) GXT4500P	IBM POWER 275 (1w) GXT6500P IBM POWER 275 (1w) GXT4500P	HP c8000 (2w) FireGL X3
Dynamic Graphics Throughput	HP c8000 (2w) FireGL X3	IBM POWER 285 (2w) GXT4500P	IBM POWER 285 (2w) GXT6500P IBM POWER 285 (1w) GXT6500P

DMU Navigator Summary Results			
Task/Test	First	Second	Third
DMU Throughput	IBM POWER 285 (2w) GXT6500P	IBM POWER 285 (2w) GXT4500P	IBM POWER 285 (1w) GXT6500P
Graphics Throughput	IBM POWER 285 (2w) GXT6500P IBM POWER 285 (1w) GXT6500P	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (1w) GXT4500P	IBM POWER 275 (1w) GXT6500P
CGR Creation Throughput	IBM POWER 285 (1w) GXT6500P IBM POWER 285 (1w) GXT4500P IBM POWER 285 (2w) GXT6500P IBM POWER 285 (2w) GXT4500P	IBM POWER 275 (1w) GXT6500P IBM POWER 275 (1w) GXT4500P HP c8000 (2w) FireGL X3	Sun Blade 2500 (1w) XVR-1200
DMU Load Throughput	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (2w) GXT6500P IBM POWER 285 (1w) GXT6500P IBM POWER 285 (1w) GXT4500P	IBM POWER 275 (1w) GXT6500P HP c8000 (2w) FireGL X3 IBM POWER 275 (1w) GXT4500P	Sun Blade 2500 (1w) XVR-1200
DMU Clash Throughput	IBM POWER 285 (2w) GXT6500P IBM POWER 285 (2w) GXT4500P	IBM POWER 285 (1w) GXT6500P IBM POWER 285 (1w) GXT4500P	HP c8000 (2w) FireGL X3

Individual Application Test Throughput Results:

Primary Application Test Throughput Results			
Task/Test	First	Second	Third
Modeling Solid Model Creation and Modification	IBM POWER 285 (1w) GXT4500P IBM POWER 285 (1w) GXT6500P IBM POWER 285 (2w) GXT6500P IBM POWER 285 (2w) GXT4500P	IBM POWER 275 (1w) GXT4500P IBM POWER 275 (1w) GXT6500P	HP c8000 (2w) FireGL X3
Finite Element Analysis (ANSOLID)	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (2w) GXT6500P IBM POWER 285 (1w) GXT4500P IBM POWER 285 (1w) GXT6500P	IBM POWER 275 (1w) GXT4500P IBM POWER 275 (1w) GXT6500P	HP c8000 (2w) FireGL X3
NC Operations STL Generation	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (2w) GXT6500P IBM POWER 285 (1w) GXT4500P IBM POWER 285 (1w) GXT6500P	IBM POWER 275 (1w) GXT6500P IBM POWER 275 (1w) GXT4500P	HP c8000 (2w) FireGL X3
Detail Drawing Creation	IBM POWER 285 (1w) GXT6500P IBM POWER 285 (1w) GXT4500P IBM POWER 285 (2w) GXT6500P IBM POWER 285 (2w) GXT4500P	IBM POWER 275 (1w) GXT4500P IBM POWER 275 (1w) GXT6500P	HP c8000 (2w) FireGL X3

Secondary Application Test Throughput Results			
Task/Test	First	Second	Third
Solid and Surface Analysis Function	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (1w) GXT4500P IBM POWER 285 (1w) GXT6500P IBM POWER 285 (2w) GXT6500P	IBM POWER 275 (1w) GXT6500P IBM POWER 275 (1w) GXT4500P	HP c8000 (2w) FireGL X3
Model Save	IBM POWER 285 (1w) GXT6500P	IBM POWER 285 (1w) GXT4500P	IBM POWER 285 (1w) GXT6500P Sun Blade 2500 (1w) XVR-1200
Walk Through	HP c8000 (2w) FireGL X3	IBM POWER 285 (2w) GXT4500P	IBM POWER 285 (1w) GXT4500P
Bend (Sheet Metal Part Development and Modification)	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (1w) GXT4500P IBM POWER 285 (2w) GXT6500P IBM POWER 285 (1w) GXT6500P	IBM POWER 275 (1w) GXT4500P IBM POWER 275 (1w) GXT6500P	HP c8000 (2w) FireGL X3
Fitting Simulation	IBM POWER 285 (2w) GXT4500P	IBM POWER 285 (1w) GXT6500P IBM POWER 285 (1w) GXT4500P IBM POWER 285 (2w) GXT6500P	IBM POWER 275 (1w) GXT6500P
Studio	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (2w) GXT6500P	IBM POWER 285 (1w) GXT4500P IBM POWER 285 (1w) GXT6500P	IBM POWER 275 (1w) GXT4500P IBM POWER 275 (1w) GXT6500P
Image Viewer	IBM POWER 285 (2w) GXT4500P IBM POWER 285 (2w) GXT6500P	IBM POWER 285 (1w) GXT6500P IBM POWER 285 (1w) GXT4500P	IBM POWER 275 (1w) GXT6500P IBM POWER 275 (1w) GXT4500P

Why TAGITT?

The Need for Application-Level Testing

When mechanical engineers select and use workstations, performance considerations should be based on the ability of the machine to rapidly complete the users' design task. Users are concerned with throughput: how much faster (or better) could I design my next product if I upgraded to a faster graphics card or a faster CPU? In the workstation industry, MHz, MIPS, MFLOPS, SPECmarks, etc. have become the standards for performance comparison of CPUs. For graphics, 3D vector drawing speed and polygon drawing speed (polygons per second) are often used for comparison. In selecting a workstation for a mechanical design application, the user is faced with a choice between many competitive machines — some with higher MIPS ratings, others with higher vector and/or polygon rates. Without running an actual application benchmark, it is difficult to predict which of the two machines will provide the better performance level for its application.

MHz, MIPS, MFLOPS, vectors per second, GPC, XPC, OPC and polygons per second all allow users to compare machines, but those specs may be misleading as predictors of engineering task efficiency. Today's CAD/CAM applications are typically very large, complicated programs. The way in which these programs perform in the context of different hardware architectures and with different operating system services and graphic libraries is generally not predictable from the previously mentioned specifications. Although software vendors are striving to make their code highly "portable" so that it runs on a wide variety of machines, the fact is that all applications must be ported and tuned to obtain optimal performance. Each workstation vendor offers unique performance-enhancing capability. Without tuning, application software may or may not take full advantage of the target hardware/operating system platform. Since software developers cannot possibly take advantage of every function in every workstation and/or operating system, performance compromises occur. The user has no way of knowing to what extent his/her application software has been ported and tuned to match capabilities offered by any particular workstation vendor without application level testing.

Test Description

TAGITT, The Albert Group Interactive Throughput Test, was designed to directly measure performance in completing typical engineering design tasks, especially related to solid modeling. For users of solid modeling software, the results provide a comparison of workstations that is more relevant than the typical manufacturer's published specifications of MHz, MIPS, MFLOPS, vectors per second and polygons per second.

TAGITT testing is typically accomplished by recording and playing back user interaction scenarios. Most CAD/CAM applications include functions to accomplish this task although some are undocumented. Record and playback mechanisms are the preferred method of testing for a number of reasons, including repeatability, accuracy, and user relevance. Although it is often easier to measure times for individual operations or functions, this can be a misleading measure of performance from a user perspective. Users constantly switch between functions and/or modules, which can result in significant performance variation as portions of the software are loaded, unloaded, and accessed from memory. The use of interaction scenarios provides a more realistic measurement of overall system performance. TAGITT tests use built-in timing and data capture mechanisms in order to obtain accurate measurements over a relatively large number of functional tests.

In addition to overall time measurements, TAGITT scenarios normally include interim times for specific functions or operations. These can provide specific performance data for individual functions such as adding a solid feature, shading a model or generating an NC tool path. These times are also used by Albert-Battaglin Consulting Group as a rough method for isolating performance that is compute intensive, graphic intensive or I/O intensive. While it is clear that the interplay between these system aspects is too complicated to be accurately measured at the application level, the measurements can sometimes point to areas for in-depth performance profiling using specialized tools.

TAGITT interaction scenarios consist of a variety of operations, with an emphasis on parametric/variational solid modeling and associated tasks such as part visualization, kinematic and "walk through" analysis, drawing view creation from 3D models, NC tool path generation, STL output, geometric/FE analysis and high quality image rendering and display. Regardless of

the task, special emphasis is placed on tasks that are not by nature “interactive.” For example, the creation of a line segment in most systems takes place in well under a half second so that performance differences will most likely be unnoticed by users. These times are not considered in TAGITT results. In contrast, updating a solid model or regenerating a drawing layout following a dimensional change can take from many seconds to several minutes and therefore has an impact on a user’s productivity. TAGITT evaluations also measure ancillary tasks such as changing functions (through menu picks) or selecting geometry. The time taken for these operations generally ranged between 1 and 15 seconds. Albert-Battaglin Consulting Group feels that the overall responsiveness of the system is reflected in these interaction times. This “responsiveness” is the difference between systems that seem heavy and slow, compared to those that “feel” quick and light. The TAGITT measurements gather data from both of these interaction types and combine them together to create an overall throughput measurement.

Models used for TAGITT evaluations are taken from previously released customer data and no data is ever provided to workstation vendors or third parties by the Albert-Battaglin Consulting Group.

TAGITT-CATIA/ENOVIA DMU

The V4 portion of TAGITT-CATIA/ENOVIA DMU consists almost entirely of a series of CATIA “record files” which are capable of recording and playing back a series of user interactions. The record files capture a majority of user interactions including some simulations of dynamic graphic manipulations performed via the GRAPER utility function. Fifty record files are used in the 4.2.4 refresh 2 version of TAGITT/CATIA V4. The record files were either created by ABCG, adapted from standard CATIA Operator Exchange test files or were developed around Dassault Systèmes’ demo part models. Some of the models used are shown in Figure 1. These files cover many areas of the CATIA 4.2.4 R2 product including part modeling, surface intersection, drawing layout, parametric modification and updates, the skin function, finite element analysis, fitting simulation, kinematics simulation, walk through analysis, NC tool path generation, sheet metal part modeling, studio image rendering and viewing and model storage and retrieval from disk. Combined, these files represent thousands of user interactions and many hours of operator seat time.



Figure 1 – Some of the CATIA models used for TAGITT/CATIA

TAGITT/CATIA also includes CATIA Image Viewer and DMU Navigator tests. For the CATIA Image Viewer test, the time for displaying an image generated by the Studio function as recorded by the Image Information function is used.

For ENOVIA DMU tests, many different CATIA V4, V5 and CGR models are combined and rendered in wireframe, shaded, shaded with edges, hidden line, shaded with material, single light, dual light, dual light with edges, neon and neon with edges modes. These models are each rotated panned and zoomed via CATScript macros. Tests are run with and without cached “cgr” files. The combination of models used for the testing of both of these functions is large and complicated enough such that the operations are generally not interactive (i.e. less than 0.5 seconds). Some of the models used for the ENOVIA DMU test are shown in Figure 2.



Figure 2 – ENOVIA DMU Navigator Test Models

Test Weighting

While running the TAGITT benchmark, the times for individual operations and scenarios are recorded. To produce Application, Graphics and Overall throughput results, weighted sums of appropriate individual test results are used. Albert-Battaglin Consulting Group sets the weighting factors for each type of operation based on its judgment of the relative importance of each operation. The weighting is Albert-Battaglin Consulting Group's best judgment for a "typical" CATIA/DMU user, whether from aerospace, automotive or any other industry. The weights are applied to actual times by averaging the results measured across the various workstations and applying the appropriate factor.

The results presented in this report represent a cross section of different types and sizes of models that can act as a guide for overall workstation performance.

Overall V4 Throughput numbers consist of the weighted sum of the Application, Graphics and Miscellaneous portions of the test. Application tests measure the times required to complete application related tasks such as changing a solid model, generating a multi-view drawing, making a STL model or running a FEM analysis. Graphics portions of the test measure exclusively viewing-related functions such as generating a shaded image or dynamically rotating that image. Miscellaneous tests include the times needed to start and exit CATIA, change CATIA functions or options, save models, and select geometry. The weighting used is shown in

Figure 3.




Figure 3 – V4 Overall Throughput Weighting

The Application time for this TAGITT test is the weighted sum of the Solid and Surface Analysis, Sheet Metal Part Development, Detail View Generation, Finite Element Analysis (ANSOLID), Fitting Simulation, Kinematics Simulation, Solid Modeling Creation and Modification, NC STL Generation, Studio, Image Viewer, and Walk Through portions of the test. The weightings used are shown in Figure 4.




Figure 4 – V4 Application Weighting Factors

V4 Graphics throughput is calculated based on the weighted times for the various graphics tests (refer to Figure 5) so as to present one representative number for this aspect of each workstation. These operations include not only the often measured dynamic graphic manipulation (dial turning) functions, but also the “graphics compute” functions which often occur the first time one accesses these operations on a given part. This combined measurement gives a better overall picture of graphic performance during typical work sessions from a user perspective. Isolated evaluations of dynamic shading or hidden line processing may be good for tuning tasks, but they do not adequately take into account the mix of operations encountered by a user. These graphic compute operations are always much longer than the dynamic manipulations themselves and are also dependent on CPU performance and its interactions with the graphics processor. The V4 Graphics Throughput Time for this TAGITT test is the weighted sum of the Graphics Compute, Hidden Line, Shaded Image and Wireframe portions of the test which are each sums from the various (10) models used throughout the testing.




Figure 5 – V4 Graphics Weighting Factors

It should be noted that user perceptible differences in display performance would, in general, only be found when processing large complex models. Overall workstation and graphic

performance across the industry has progressed to where “simpler” models can be easily transformed interactively. Unfortunately, we found no easy method for defining a “simple” model. Model size and the number of geometric elements (surfaces, planes, lines, etc.) did not correlate with graphic performance. The performance is based on the complexity of the various geometric components, which is not easy for the user to determine. It is difficult to determine what class of workstation offers “sufficient” performance without examining explicitly the types of models and displays commonly used.

V4 Responsiveness is calculated based on the weighted sum of the CATIA start and exit times, the times to change options or “load” V4 function, and the time required to select geometry in very large models. As with all TAGITT measurements, the emphasis is placed on interactions which generally take more than one second and can take up to thirty seconds on some systems. The weightings used are shown in Figure 6.





Figure 6 – V4 Responsiveness Weighting Factors

ENOVIA DMU Navigator throughput is calculated based on the weighted sum for the Graphics, Load, Interference Check and CGR Create tests. The weighting factors are shown in Figure 7.



.Figure 7 – ENOVIA DMU Navigator Weighting Factors

ENOVIA DMU Navigator Graphics weighting factors are shown in Figure 8.




Figure 8 – ENOVIA DMU Graphics Navigator Weighting Factors

Machines Evaluated

The following table shows the machines tested and their configurations. All machines were configured with 2GB main memory, 5GB swap, 19”/21” color monitor, CD/DVD ROM drive, Ethernet interface, mouse, keyboard and 100GB or larger hard disk, operating system, and 3D API. Note: throughout the report, references to the HP and Sun machines include the graphics accelerators listed below.

Vendor	Model	Graphics	CPU	L2 Cache MB	CPU MHz	OS	Abbreviation used in report and on charts
IBM	IntelliStation POWER 285	GXT6500P	POWER5+™ (2w)	1.9*	1900	AIX 5L™ v5.3	IBM POWER 285 (2w) GXT6500P
IBM	IntelliStation POWER 285	GXT4500P	POWER5+ (2w)	1.9*	1900	AIX 5L v5.3	IBM POWER 285 (2w) GXT4500P
IBM	IntelliStation POWER 285	GXT6500P	POWER5+	1.9*	1900	AIX 5L v5.3	IBM POWER 285 (1w) GXT6500P
IBM	IntelliStation POWER 285	GXT4500P	POWER5+	1.9*	1900	AIX 5L v5.3	IBM POWER 285 (1w) GXT4500P
IBM	IntelliStation POWER 275	GXT6500P	POWER4+™	1.5**	1450	AIX 5L v5.3	IBM POWER 275 (1.45 GHz) GXT6500P
IBM	IntelliStation POWER 275	GXT4500P	POWER4+	1.5**	1450	AIX 5L v5.3	IBM POWER 275 (1.45 GHz) GXT4500P
HP	c8000	ATI FireGL X3	PA-8800 (2w)	1.5(L1) 32(L2)	1000	HP-UX 11i	HP c8000 (2w 1.0 GHz) Fire GL-X3
Sun	Blade 2500	XVR-1200	UltraSparc-IIIi	1	1600	Solaris 5.10	Sun Blade 2500 (1.6 GHz) XVR1200

*Also has 36MB of L3 cache

**Also has 8MB of L3 cache

NOTE FOR ALL PERFORMANCE CHARTS: All performance charts in this report show the throughput for each system compared versus the slowest system for each test. For this reason longer bars mean better performance. A bar value of 100% means that the given system is twice as fast as the slowest system. The numbers to the right of the bars represent throughput time. For these values lower numbers represent better performance.

TAGITT-CATIA/ENOVIA DMU Results

Overall V4 Throughput

Chart 1 compares the overall weighted throughput for the CATIA V4 interactive scenarios including the graphics tests. Albert-Battaglin Consulting Group feels that this number gives the best overall rating of workstation performance running CATIA V4. The chart shows the IBM POWER 285 (2w) GXT4500P to be the fastest machine overall followed closely by the IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (1w) GXT6500P. The third place IBM POWER 285 (1w) GXT4500P machine was 7% slower than the IBM POWER 285 (2w) GXT4500P in overall throughput. The HP c8000 (2w) FireGL X3 was the fastest machine from a competitor, a substantial 70% slower than the leader. The Sun Blade 2500 (1w) XVR-1200 was a disappointing 139% slower than the leader.




Chart 1 – Overall V4 Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

V4 Graphics Throughput

Chart 2 shows the weighted cumulative time to complete V4 graphic view manipulation operations of parts and assemblies. As described earlier, the Albert-Battaglin Consulting Group Graphics Throughput time includes both initial “loading” of graphics as well as dynamic manipulation times. The HP c8000 (2w) FireGL X3 and IBM POWER 285 (2w) GXT4500P were the fastest machines in this test with a time 9% faster than the second place IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (1w) GXT6500P machines. In third place, the IBM POWER 285 (1w) GXT4500P was 21% slower than the leader. The Sun Blade 2500 (1w) XVR-1200 finished last in this test, 85% slower than the leader.

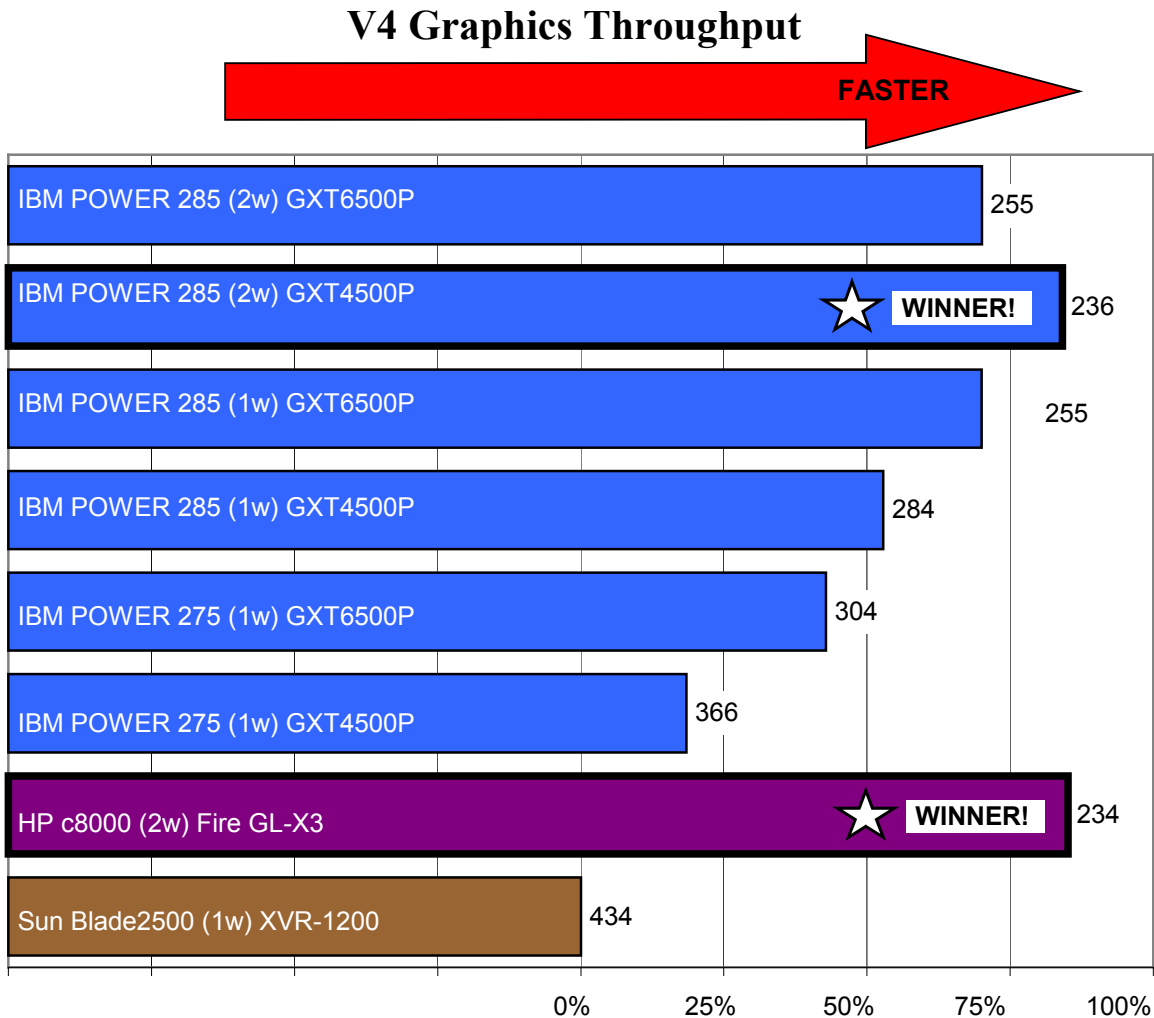


Chart 2 – V4 Graphics Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

V4 Application Throughput

Chart 3 shows the weighted cumulative time to complete all of the CPU intensive application-specific tasks in the V4 benchmark such as solid modeling operations, drafting and detailing operations, FEM functions and NC computations. All four IBM POWER 285 workstations were essentially tied for first place. In a two-way tie for second place, the IBM POWER 275 (1w) GXT6500P and IBM POWER 275 (1w) GXT4500P workstations were 46% slower than the leader. The third place HP c8000 (2w) FireGL X3 machine finished 90% slower than the leader, while the Sun Blade 2500 (1w) XVR-1200 was finished 160% slower than the leader.

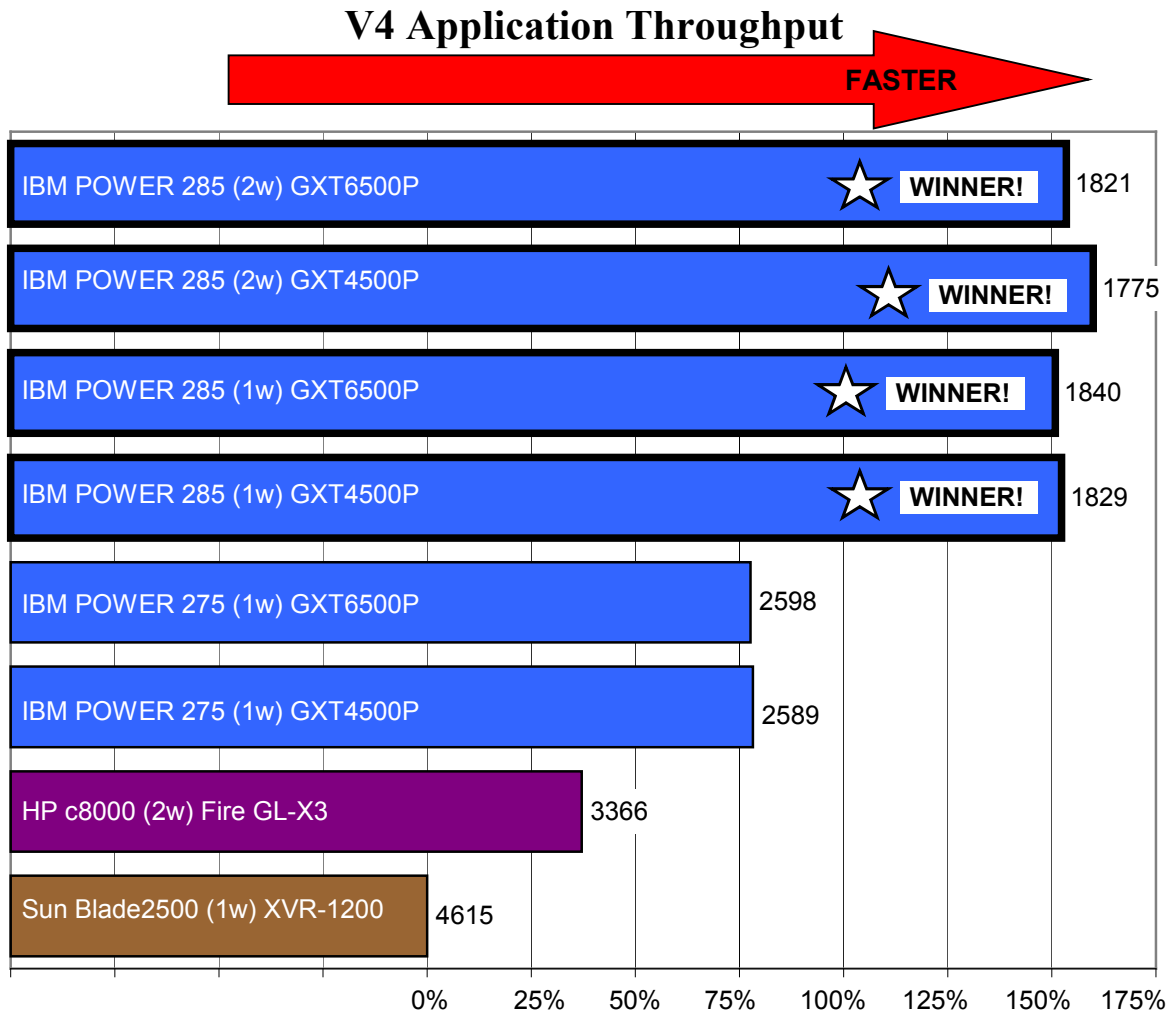


Chart 3 – V4 Application Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

V4 System Responsiveness

Chart 4 shows the cumulative time to complete all of the V4 system responsiveness tests in the benchmark. The tests measure the “quickness” of the system, doing common CATIA V4 tasks such as starting and exiting the application, changing functions and selecting elements. Again all four IBM POWER 285 workstations were the overall winners of this test. Tied for second place, the IBM POWER 275 (1w) GXT4500P and IBM POWER 275 (1w) GXT6500P machines were 33% and 35% slower than the leading machines. In third place, the Sun Blade 2500 (1w) XVR-1200 was 50% slower than the fastest machines. In this set of tests, the HP c8000 (2w) FireGL X3 finished last, just over twice as slow as the leaders.




Chart 4 – V4 System Responsiveness Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

V4 Dynamic Graphic Throughput

shows the comparative throughput performance of the V4 dynamic graphics tasks in the benchmark including wireframe, shaded and hidden line dynamic graphic operations. These results do not include the longer graphic computation times that are included in the overall graphics throughput results. The winner of this test was the HP c8000 (2w) FireGL X3, finishing a strong 29% faster than the IBM POWER 285 (2w) GXT4500P in second place. The IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (1w) GXT6500P machines were 41% slower than the leading machine in third place. The Sun Blade 2500 (1w) XVR-1200 finished a disappointing 114% slower than the leader.

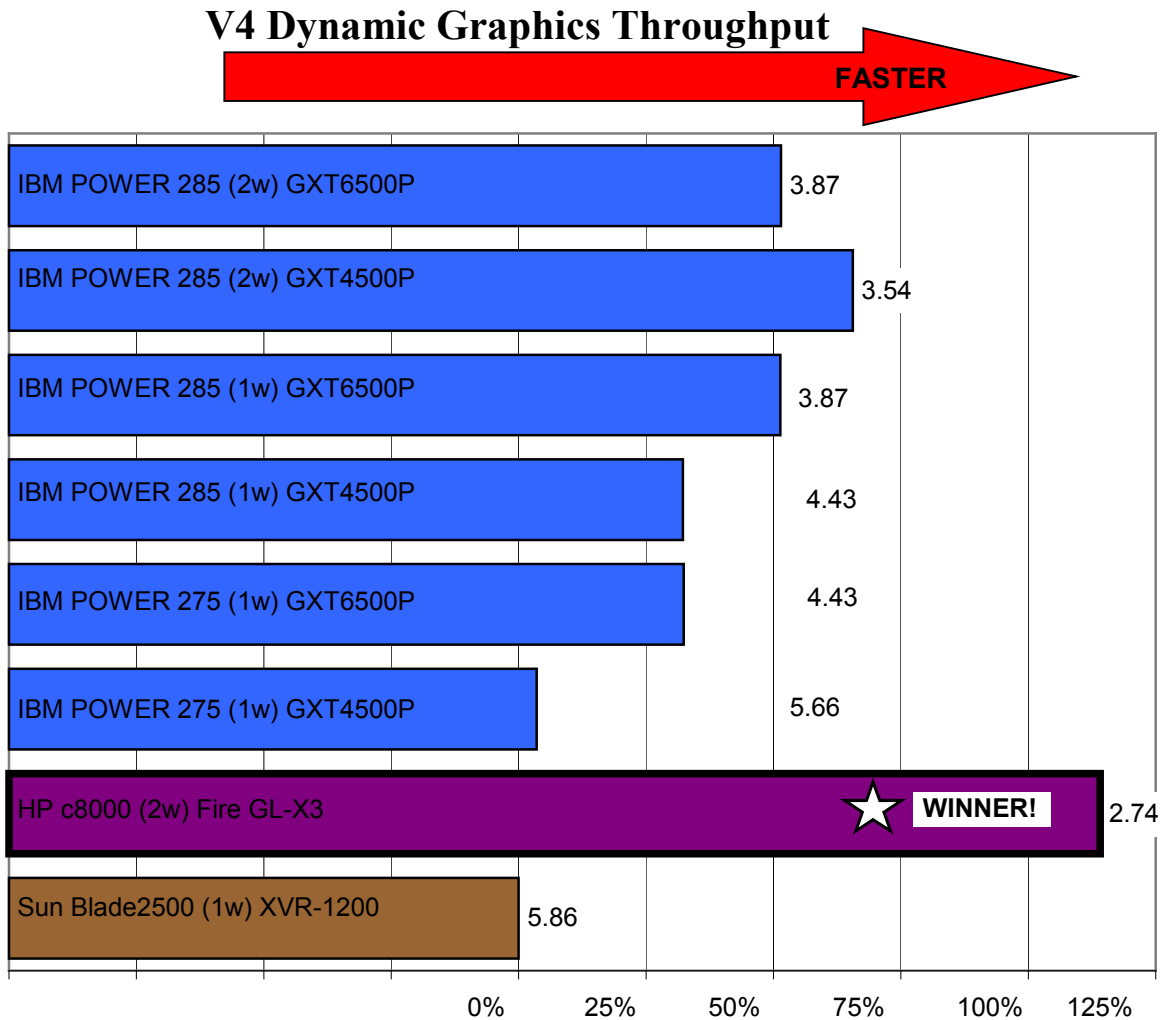


Chart 5 – V4 Dynamic Graphic Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

ENOVIA DMU Navigator Overall Throughput

Chart 6 shows the weighted cumulative time to complete all of the ENOVIA DMU Navigator tests in the benchmark. The ENOVIA DMU Navigator uses OpenGL graphics and the TAGITT tests work with V4, V5 and “cgr” data so these results are important for users working with both CATIA V4 and CATIA V5. The overall winners of this test were the IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (1w) GXT6500P. In second place, the IBM POWER 285 (2w) GXT4500P and IBM POWER 285 (2w) GXT6500P machines were 13% and 17% slower than the leader. The third place IBM POWER 275 (1w) GXT6500P machine was 56% slower than the leader. The HP c8000 (2w) FireGL X3 finished the test 58% slower than the leaders while the Sun Blade 2500 (1w) XVR-1200 lagged behind 124% slower than the fastest machines.




Chart 6 – ENOVIA DMU Navigator Overall Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

ENOVIA DMU Navigator Graphics Throughput

Chart 7 shows the comparative throughput for the ENOVIA DMU Navigator graphic tests in the benchmark. Because the DMU navigator is so often used for exploring and checking designs visually, this capability was highly weighted in the testing. As with all TAGITT graphics test, these results concern themselves with very complicated scenes involving between 1.9 and 22.9 million triangles. With these complex scenes all systems are tested to their limits. The overall winners of this test were the IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (1w) GXT6500P. In second place, the IBM POWER 285 (2w) GXT4500P and IBM POWER 285 (2w) GXT6500P machines were just over 30% slower than the leader. The third place IBM POWER 275 (1w) GXT6500P machine was 52% slower than the leader. The HP c8000 (2w) FireGL X3 finished the test 58% slower than the leaders while the Sun Blade 2500 (1w) XVR-1200 lagged behind 133% slower than the fastest machines.




Chart 7 – ENOVIA DMU Navigator Graphics Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

ENOVIA DMU Navigator Interference Check Throughput

Chart 8 shows the comparative throughput for the ENOVIA DMU Navigator Interference Computation tests in the benchmark. Because these operations can be very time consuming, this capability was highly weighted in the testing. The overall winners of this test were the IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (2w) GXT4500P. In second place, the IBM POWER 285 (1w) GXT6500P and IBM POWER 285 (1w) GXT4500P machines were over 77% slower than the leader a clear advantage for the multiprocessor configurations. The HP c8000 (2w) FireGL X3 machine was 136% slower than the leader despite being a two-way system. The Sun Blade 2500 (1w) XVR-1200 was incredibly slow in this particular test, taking 6 times longer to finish.




Chart 8 – ENOVIA DMU Navigator Interference Check Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

Conclusions

The results of the TAGITT-CATIA V4/ENOVIA DMU evaluation show IBM to have a solid overall performance lead with its new 1.90 GHz POWER 285 machines based on the POWER5+ architecture. In V4 tests, the IBM POWER 285 (2w) GXT4500P was 70% faster overall than the nearest competitive machine, the HP c8000 (2w) FireGL X3. The Sun Blade 2500 (1w) XVR-1200 was no match for the IBM machines with overall performance times a disappointing 139% slower than the leader. In ENOVIA DMU testing, the IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (1w) GXT6500P were the fastest machines overall. The closest competitive machine the HP c8000 (2w) FireGL X3 was 58% slower than the leader. The Sun Blade 2500 (1w) XVR-1200 was considerably slower than the other machines tested, finishing the ENOVIA DMU tests 124% slower than the leader.

The new IBM POWER 285 workstations were overall about 40% faster than 1.45GHz IBM POWER 275 workstations. In the ENOVIA DMU testing, the difference was more than 50%.

In terms of overall V4 graphic performance, the HP c8000 (2w) FireGL X3 tied with the IBM POWER 285 (2w) GXT4500P as the fastest performer. This shows the strength of IBM's graphic strategy allowing the faster POWER5+ processor to drive excellent graphic performance with the proven GXT4500P graphics adaptor. The fact that the GXT4500P outperformed the GXT6500P on the 2-way system is attributed to IBM's multi-threaded graPHIGS API implementation and the extreme speed of the POWER5+ processor to do the graphics pipeline processing for the GXT4500P. For pure dynamic graphic response, the HP c8000 (2w) FireGL X3 outperformed the IBM POWER 285 (2w) GXT4500P by 29%. Sun's dynamic graphic results were notably poor; the Sun Blade 2500 (1w) XVR-1200 finished 114% slower than the fastest HP box.

In terms of overall ENOVIA DMU Graphic performance, the full OpenGL hardware acceleration of the GXT6500P combined with the fast POWER5+ processor(s) proved a winning combination. The IBM POWER 285 (2w) GXT6500P and IBM POWER 285 (1w) GXT6500P outperformed the second place IBM POWER 285 (2w) GXT4500P and IBM POWER 285 (1w) GXT4500P machines by just over 30%. The IBM POWER 275 (1w) GXT6500P machine was 52% slower than the leaders, while the HP c8000 (2w) FireGL X3 tested 58% slower.

The wide range of application testing in TAGITT showed performance differences between various CATIA functions. It is impossible to tell from the TAGITT evaluation the reasons for these differences. Although the 1.90 GHz IBM POWER 285 machines won nearly all of the individual application tests, there were exceptions. As stated before, the HP c8000 (2w) FireGL X3 won the dynamic graphics test as well as the Walk Through tests. In the Geometric Analysis tests on the other hand, the Sun Blade 2500 (1w) XVR-1200 machine outperformed the HP c8000 (2w) FireGL X3.

Although the HP c8000 machines tested had 2 processors, no clear advantage was apparent in this configuration when running CATIA V4 or ENOVIA DMU Navigator. Since CATIA V4 and ENOVIA DMU is mostly single threaded code, we have typically found very little advantage to multi-processor configurations. In fact, the c8000 machines did not outperform the single processor IBM POWER 275 machines in the Studio Viewer tests, which in the past have benefited from multiprocessor technology.

The dual processor IBM POWER 285 configurations did show performance advantages in several tests. The overall win of the IBM POWER 285 (2w) GXT4500P in the V4 testing shows a dual processor advantage as the uniprocessor configuration finished 7% slower. Individual tests showed significant performance benefits. In the Studio tests, which generate photo realistic images using a ray-tracing algorithm, the dual processor POWER 285s were about 60% faster than the single processor models. In the Image view application, the dual processor 285s were almost twice as fast as single processor configurations. In the ENOVIA DMU Interference computations, the 2 way 285s were nearly 80% faster than 1 way configurations.

IBM has increased its overall CATIA V4 and ENOVIA DMU performance leadership with its 1.9 GHz IntelliStation POWER 285 workstations. The TAGITT-CATIA/ENOVIA DMU test shows the POWER 285 machines to be excellent all around performers in the CATIA environment. The combination of raw CPU speed and well-balanced graphic performance in the IBM IntelliStation POWER 285 workstation line, make them excellent choices in CATIA engineering environments.

Methodology

All tests were conducted by Albert-Battaglin Consulting Group personnel. Test conditions were set up to minimize any environmental differences with the various systems. Systems were tested in “lab” environments so that they were isolated from network interference. Albert-Battaglin Consulting Group saw no evidence to suggest that performance was impacted by extraneous network activity. Almost all timing data was automatically recorded and transferred directly into spreadsheets for analysis. All tests were run at least three times and the average times were used for comparison. For the overall test, the time differences between runs were typically less than 1.5% and in most cases as low as 0.1%.

All TAGITT-CATIA V4/ENOVIA DMU tests were completed using released CATIA 4.2.4 R2 and ENOVIA R15 SP 2 software. In all cases, the software and data and licenses were loaded locally on each workstation. Manufacturers’ required and recommended software “patches” or upgrades for CATIA V4 and ENOVIA V5 were installed on all systems prior to testing. IBM machines were tested with SMT turned off.

Appendix

The following charts show the results for individual application tests that were combined to form the overall application and overall throughput times for TAGITT-CATIA V4/ENOVIA DMU.

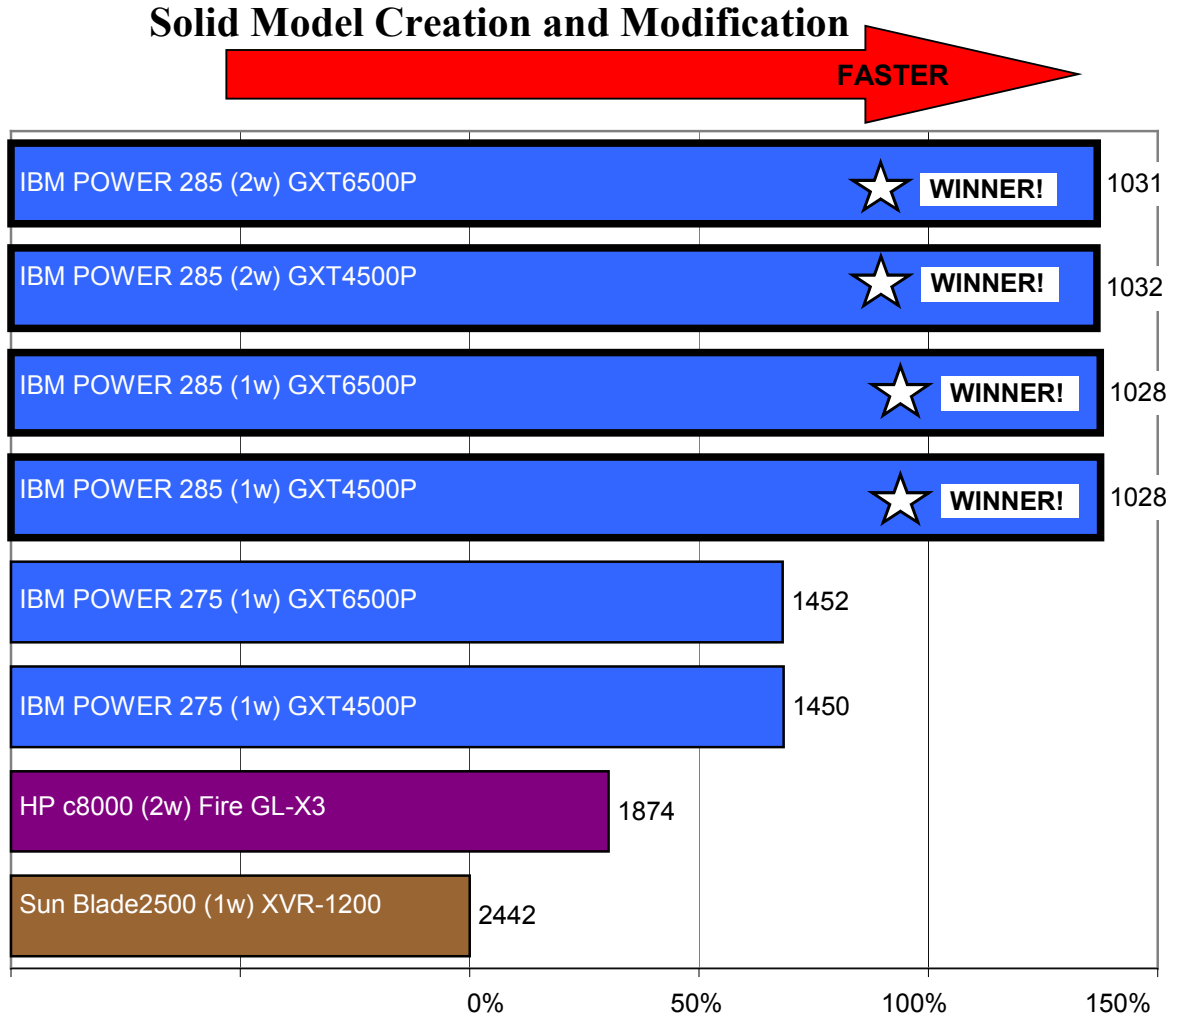


Chart 9 – Solid Model Creation and Modification Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

Finite Element Analysis (ANSOLID)




Chart 10 – Finite Element Analysis (ANSOLID) Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

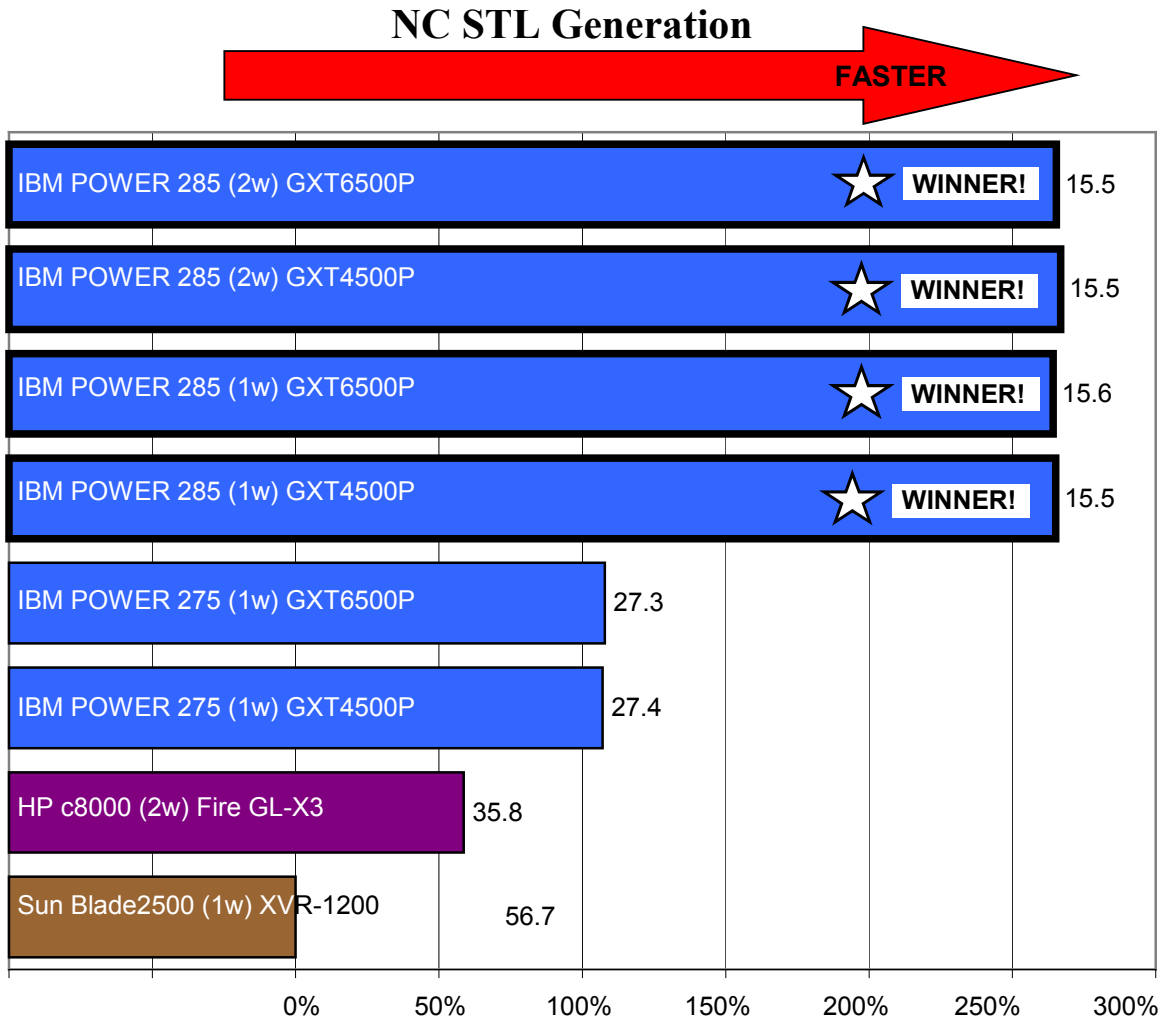


Chart 11 – NC STL Generation Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)




Chart 12 – Detail View Generation Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

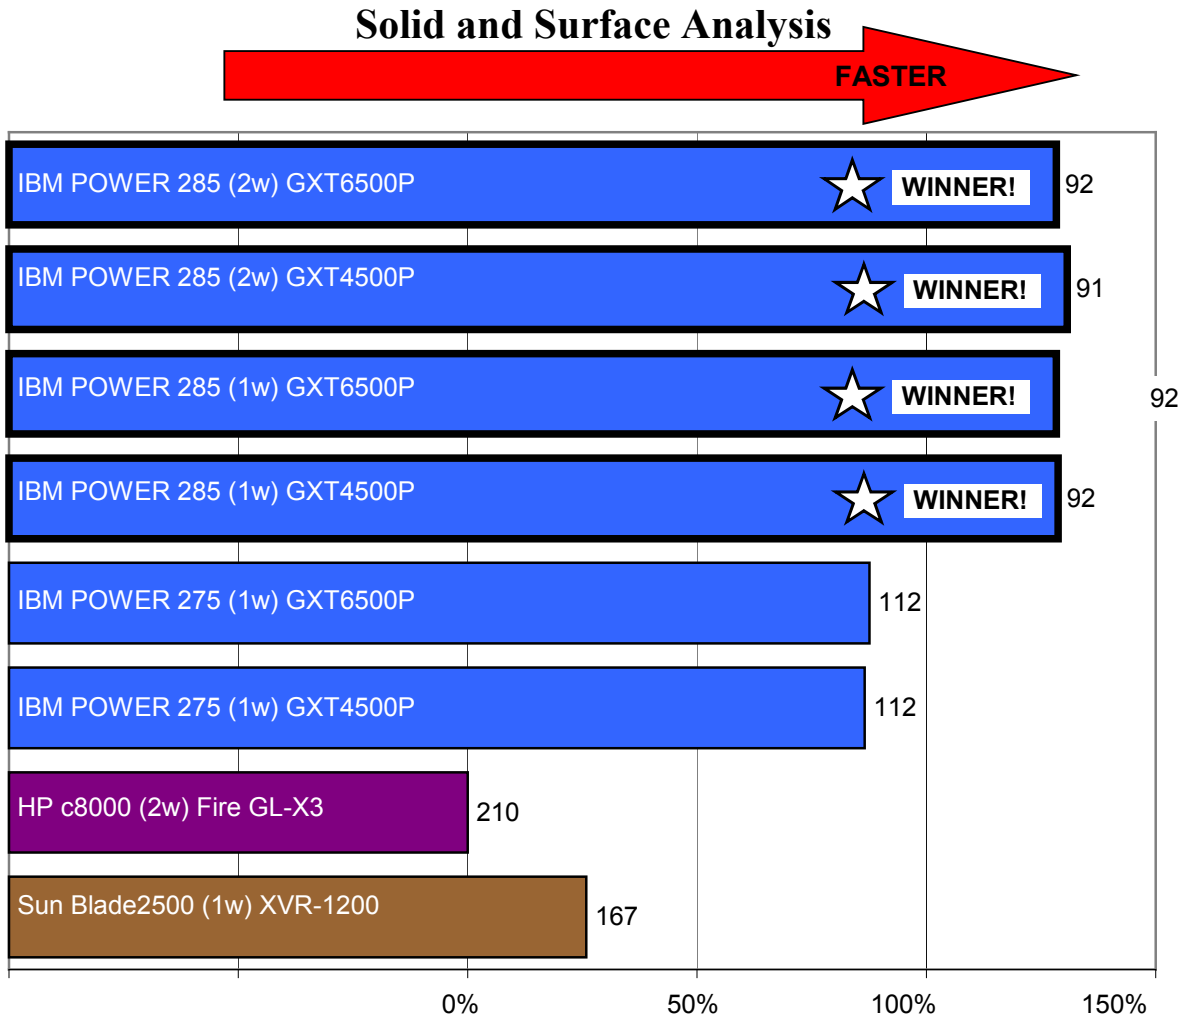


Chart 13 – Solid and Surface Analysis Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

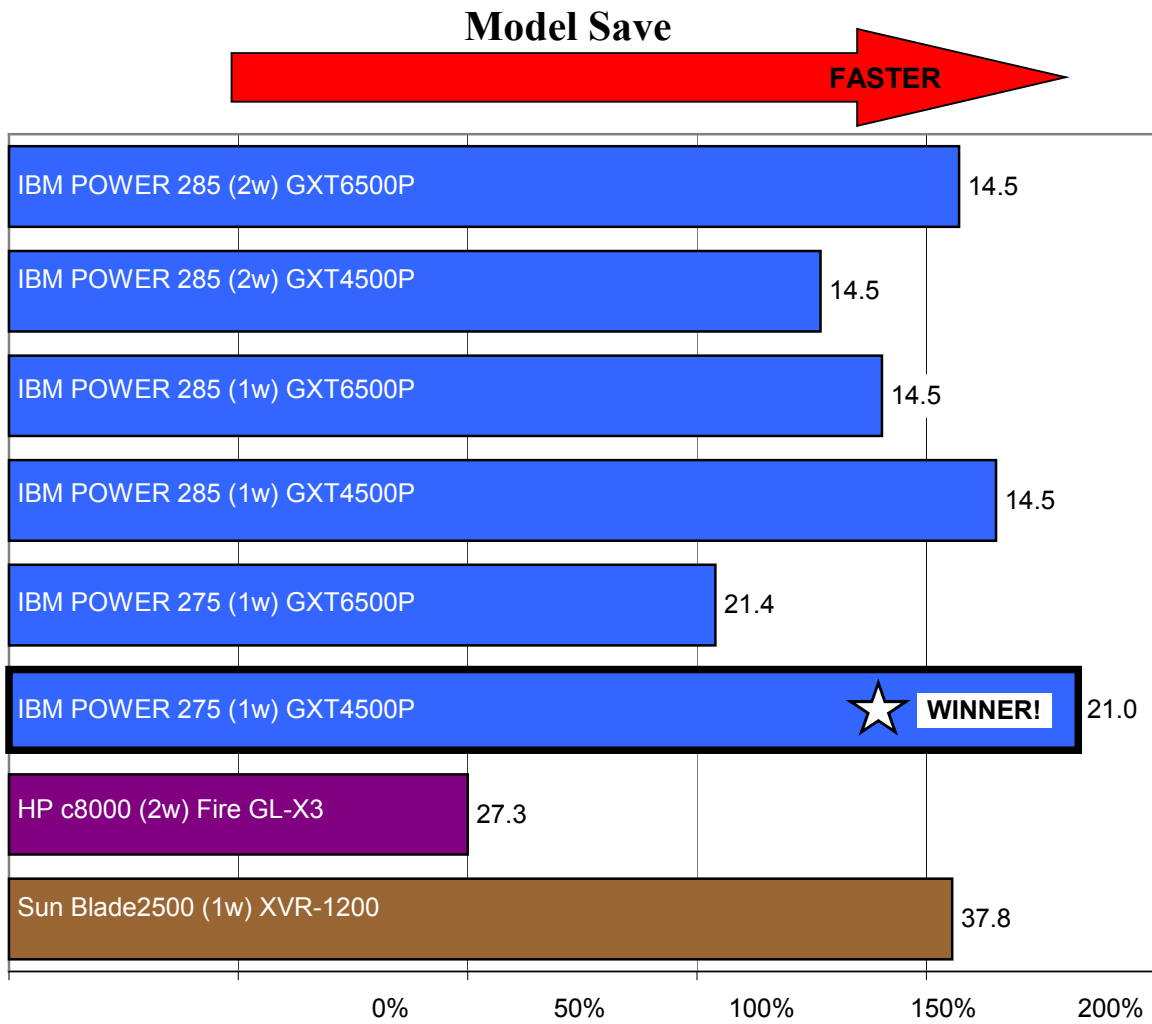


Chart 14 – Model Save Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)




Chart 15 – Walk Through Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)




Chart 16 –Sheet Metal Part Development Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

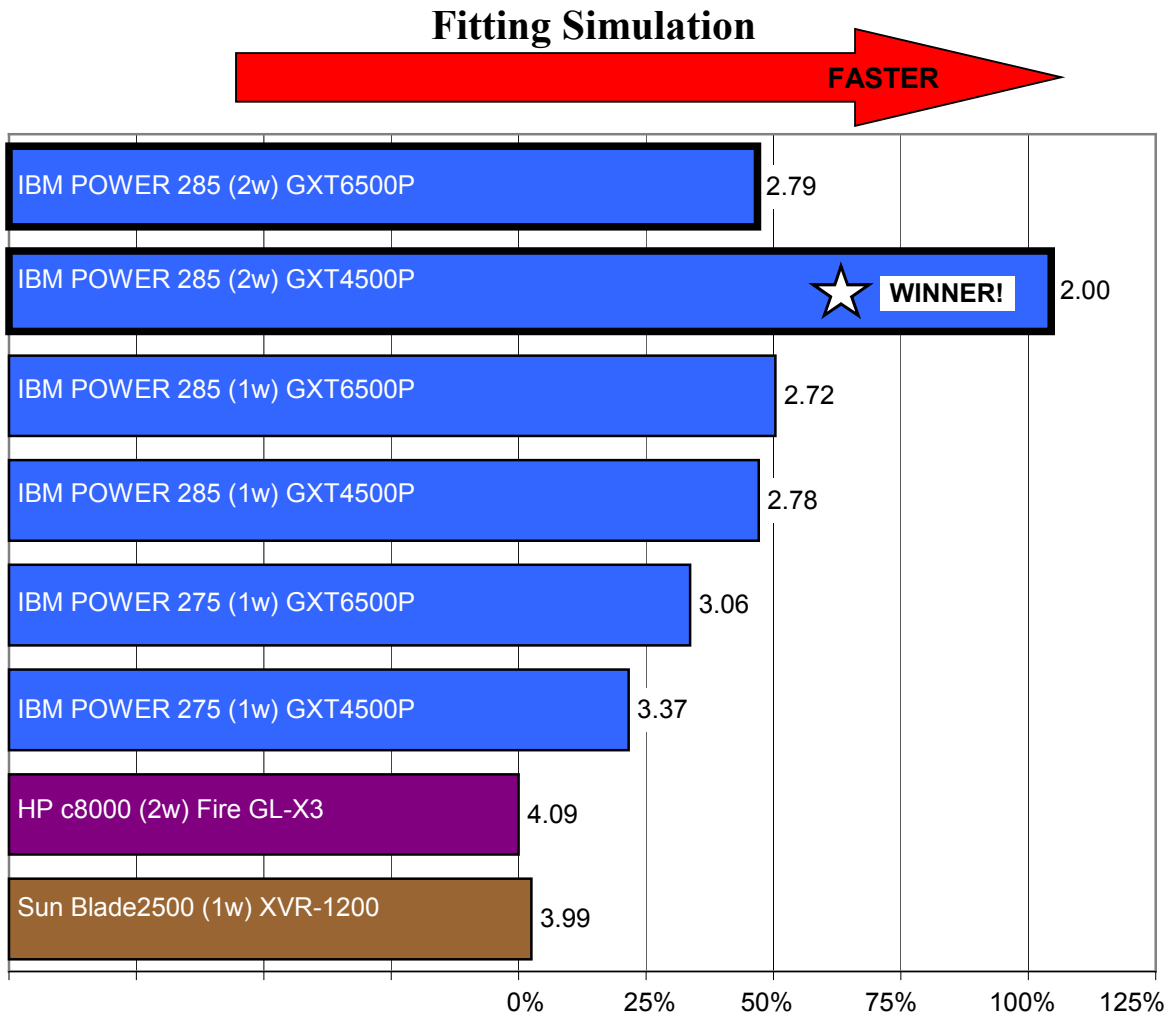


Chart 17 – Fitting Simulation Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

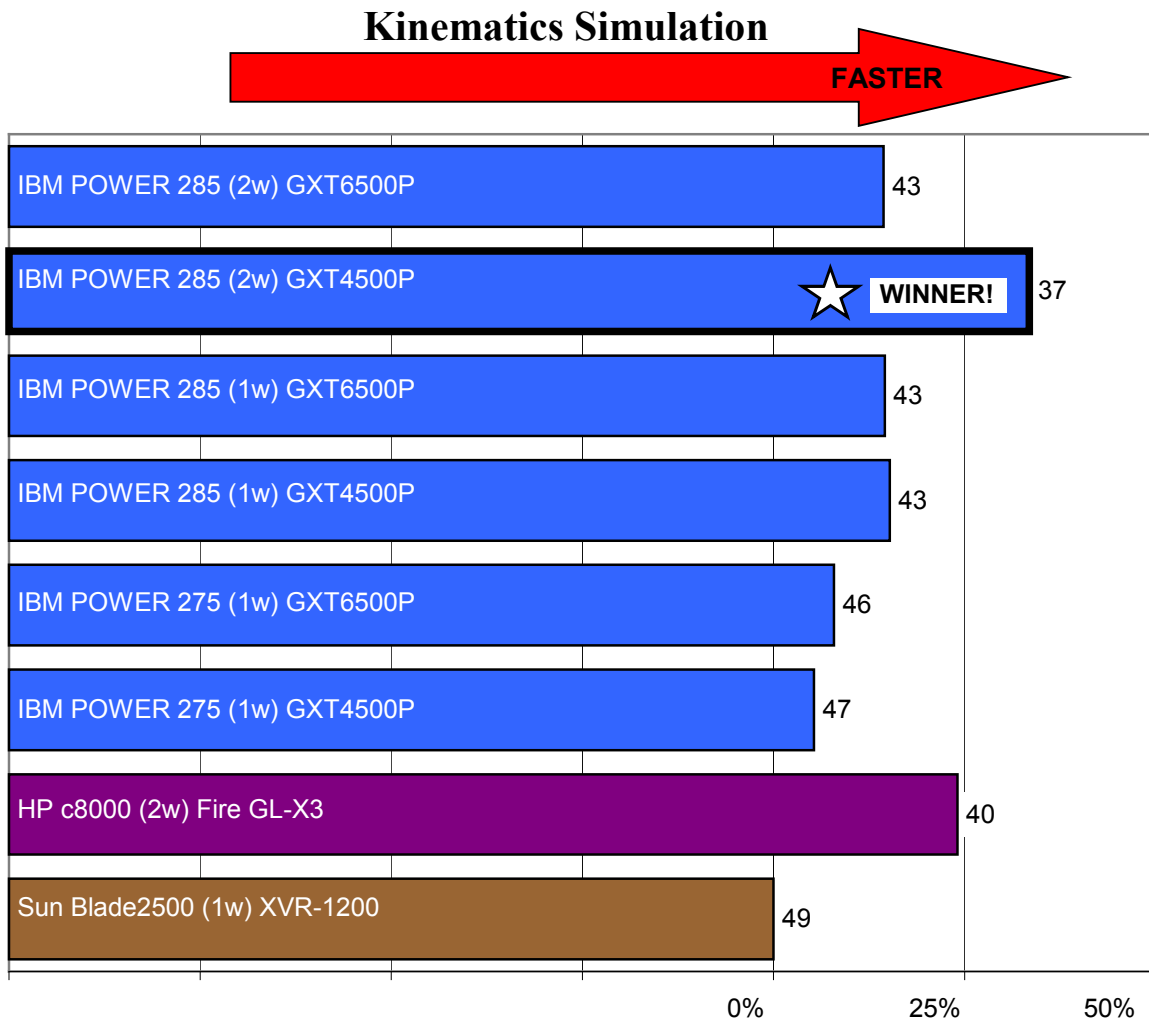


Chart 18 – Kinematics Simulation Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

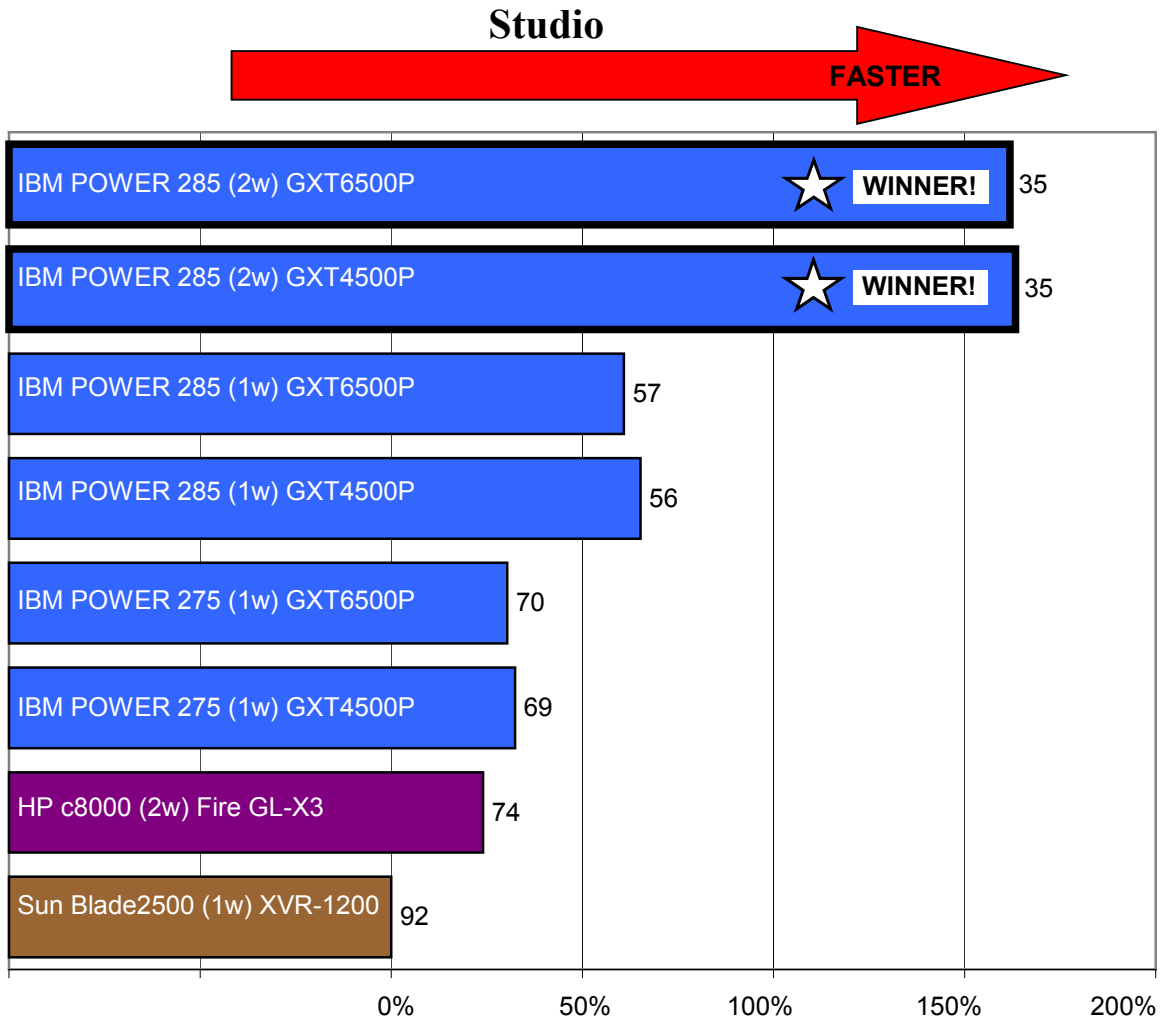


Chart 19 – Studio Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)




Chart 20 – Image Viewer Throughput Relative to Slowest Machine
Test time in seconds shown next to bars (smaller numbers faster)
Longest bar wins! (★)

ENOVIA DMU Navigator CGR Creation




Chart 21 – ENOVIA DMU Navigator CGR Creation Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)

ENOVIA DMU Navigator Load

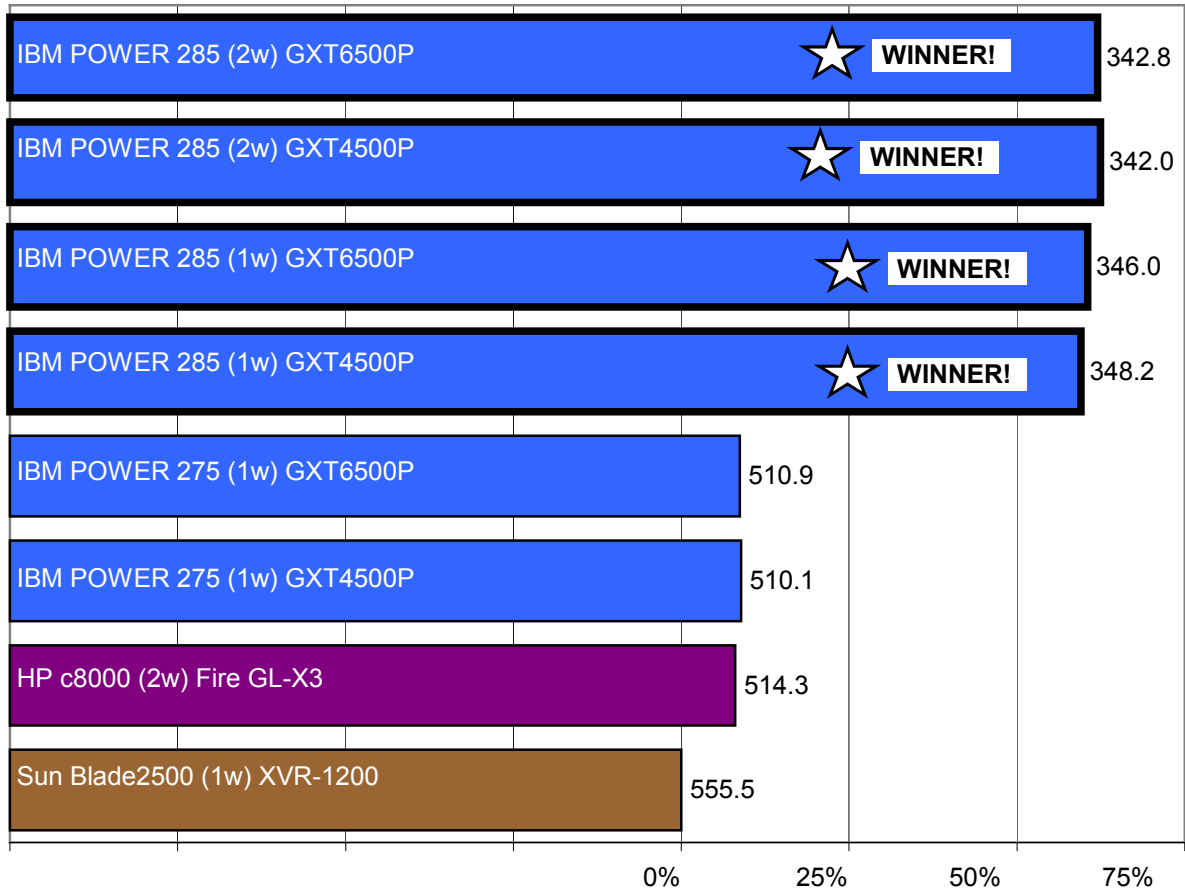



Chart 22 – ENOVIA DMU Navigator Load Throughput Relative to Slowest Machine
 Test time in seconds shown next to bars (smaller numbers faster)
 Longest bar wins! (★)