

Grids and high performance computing

Optimising grids with analytics

the sheer scale of infrastructure and volume of work involved makes grid optimisation a complex, costly and risk-laden area”



Introduction

Since its origins in academic and scientific institutions in the 1990s, high performance (HPC) grid computing has fast established itself as an enterprise architecture for reliably processing massive, complex computations at speed. Investment banks - whose IT teams pioneered the use of grids for performing risk and pricing calculations in derivatives and currency trading - are now so critically dependent on them that their business' profitability and competitiveness rests on their optimisation. With so much at stake, continually improving the grid has become a ubiquitous pursuit for organisations involved in their operation.

The criticality of grids to investment banks

This being the case, there are a number of market conditions and challenges that are having a significant impact on many organisations' efforts today. Of particular consequence is the fallout that the global economic crisis has had on the investment banking sector. Merger and acquisition activity is widespread and has resulted in many organisations being faced with large-scale change programmes that may include trade portfolio mergers, consolidation or grid relocation. Similarly, increased compliance and regulatory requirements have placed further pressure on teams to provide more robust evidence on trading processes for internal audits and industry bodies. Meanwhile, the ever-present requirement to reduce latency and process trades and risk calculations at breakneck speeds has become a vitally important differentiator in today's highly competitive marketplace. Compounding it all, the sheer scale of infrastructure and volume of work involved makes grid optimisation a complex, costly and risk-laden area.

Challenges to optimisation

Central to dealing with all of the above challenges is a clear and thorough understanding of the business demands being placed on the grid, so that applications, models and tasks requiring its use can be clearly understood and prioritised. Despite this goal, many IT teams are nonetheless solely reliant on traditional infrastructure-based forecasting techniques to manage the grid and its workload. The inhibiting issue here is that these approaches tend to be over-simplified, not focussing adequately enough on understanding business demand and in particular how jobs, applications and their trades consume grid resources. For example, some IT organisations might use time series forecasting which looks at historical utilisation data to try and find a trend, but because this approach is based purely on historical infrastructure data only (month-to-month changes in IT data), it suffers from not being correlated to the business and lacks the ability to provision and forecast accurately to changes in demand.

Moreover, with the grid and its functions remaining a fundamental and critical part of the business' daily operations, any initiative involving its reconfiguration can be potentially damaging to core operations and revenue streams, and so must be approached with full understanding of the risk and impact of change.

Grid architecture

Although grid architectures (see Fig. 1) will vary depending on their use, the core processing architecture is comprised of a server pool known collectively as the grid engines. Each server is typically a multi-CPU device running a variant of a standard Windows or Linux OS. Servers are typically grouped into partitions, which are dedicated to exclusive access for particular applications, jobs and tasks - enabling the workload to be load balanced and grid utilisation maximised.

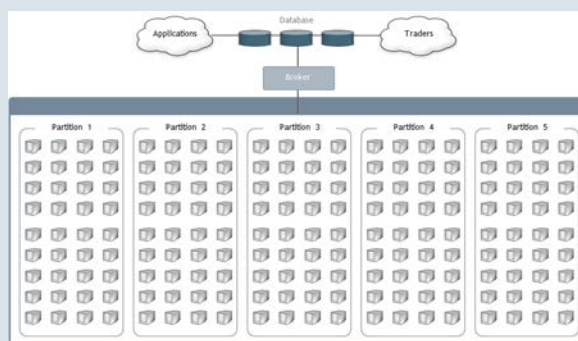


Fig. 1 – Example grid architecture

Tasks are managed by way of a broker server running grid management software, such as TIBCO/Data Synapse, Sun Grid Engine, or Platform Symphony amongst many others. The broker allocates and schedules tasks depending on their priority and relative demand from the business. For example, in a trading environment the grid may be partitioned by asset class, and jobs prioritised according to the type of analytics being done (risk, cash-flow, PnL and so on). In order to complete tasks an engine may need to interface with database and application layers to retrieve trade data and relay results back to end-users and other application components such as database and storage.

For example, when looking at ways to improve the throughput performance of a particular application using the grid, organisations must clearly understand the relationship between trades, models and engine tasks being executed. Similarly, if considering a relocation or move to an outsourced provider, careful consideration should be given to the impact of latency on grid throughput and to bandwidth requirements, so that performance will be above or on a par with the current in-house arrangement. Therefore, to fully satisfy not only cost, but risk and impact implications, IT organisations should base their grid decision making on informed, quantified facts that relate the grid's operation and performance to current and future business demand and objectives. It is with these requirements in mind that analytics is providing the essential degree of accuracy and correlation to business activity that managing today's grid environments requires.

Applying analytics to grids

By accurately modelling and analysing the complex relationships between the grid environment and application transactions, models/reports and tasks running on it (see Fig. 2 – Modelling grid data), analytics provides a “joined up” model of the current working environment, exposing hidden correlations that indicate where changes, on both the IT and business side of the equation, will bring about the greatest returns from a performance, capacity and cost perspective.

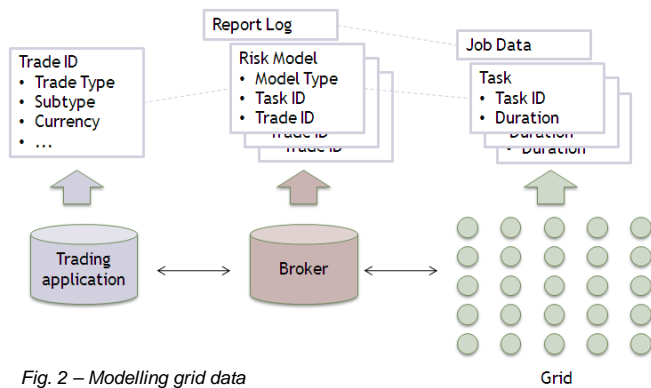


Fig. 2 – Modelling grid data

Unlike traditional approaches that only provide simplified infrastructure utilisation and forecast out from these, analytics brings together a powerful combination of techniques and expertise for truly business-aligned grid management. This advanced understanding enables both IT and the business to improve processes, de-risk change and instil proactive measures for ensuring the grid is provisioned to serve the best interests of the business and its future plans.

Baselining business volumes to the grid

The first stage of an analytics approach to optimising the grid is to create a baseline of the business load placed upon it, by capturing trade type and volume metrics that comprise the transaction portfolio. Different trades types require different analytical models to be run, and so make different demands on grid resources and overall compute time (see Fig. 3). For example, the compute effort for a Bermudan Swaption is likely to be orders of magnitude higher than for the equivalent European-style, so a knowledge of the trade portfolio is essential in predicting grid power (and hence time) required. In itself, gaining an appreciation into the mix of business transactions and models running on the grid can open up a strategic dialogue between IT and the business, by enabling an understanding into how changes to the mix can bring about improved trading performance as well as IT utilisation.

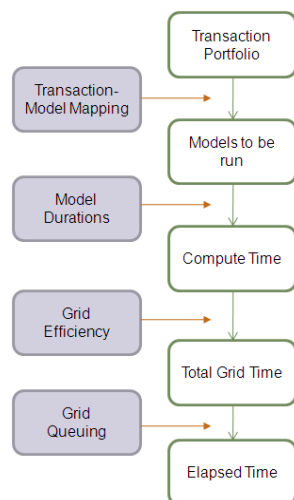


Fig. 3 – Understanding grid capacity and compute duration

Key approach and benefits

- **Baselines trading volumes and grid utilisation** – models granular trading portfolio activity to grid utilisation for precision provisioning
- **Calculates grid compute and headroom** – reveals current level of performance and capacity status with available headroom for growth
- **Optimises trading performance** – with deep-dive insight to reduce trading latency and clear bottlenecks
- **Manages and de-risks change** – keeps pace with change with active post-change monitoring and scenario modelling

The second stage creates a baseline of the grid's current utilisation, to establish the current levels of performance, utilisation and capacity. This is achieved by capturing samples of utilisation and performance metrics captured from the grid engine's underlying system components, such as node/CPU utilisation, I/O, network latency and disk storage used. This baseline is then tracked over time to ascertain the grid's utilisation by hour, day, week, and month – capturing the full range of utilisation from minimum to peak outliers, and not just averages which can distort the true picture.

Case study 1 – Aligning the grid to the business

A large UK investment bank approached Sumerian to help improve capacity planning and reduce risk within its critical trading services, which uses an HPC grid for its overnight batch risk calculation models. Sumerian's starting point was to capture and model data from the grid and its associated service components such as network management tools, SAN monitors, server metrics, message buses (such as TIBCO Rendezvous), batch log files, and trade audit trails – all of which provide rich information for understanding the grid's performance and capacity.

After constructing baseline models of the IT environment, Sumerian could then gather metrics from the business trading side, to build a picture of how different trade types correlated to the models that composed the batch run. By then applying highly granular regression analysis, Sumerian could then identify the precise relationships between the overnight batch run and utilisation of the grid – and as a result, build a predictive model for accurate capacity management and forecasting. Sumerian's findings quantified the correlation between the length of the total aggregate processing time and the daily number and mix of trades in the database. By overlaying various forecasted growth scenarios in trading activity, Sumerian could then provide the bank with precise indications on how the processing time would increase with the rise in different trade types.

The improved understanding around grid capacity and performance has enabled the bank's team to reduce risk and cope better with demand. For example, by creating a grid footprint for each trade and model type, Sumerian was able to identify those models which posed the greatest risk to overnight batch durations, enabling the development effort to focus on those models.



Understanding grid compute and headroom

The outputs from the two baselines provide the necessary elements with which to calculate the grid's current compute duration and available capacity headroom and directly relate these to business load. For example, if any areas are approaching limits, the direct correlations formed between grid utilisation and business load will determine what impact this is currently having on overall grid efficiency, resulting job queues and where bottlenecks are appearing. Similarly, the headroom calculation will clearly show whether over provisioning, or imbalanced loads are evident, enabling organisations to make evidence-based decisions to reduce capacity, consolidate or reconfigure where necessary. Conversely, where business growth is an objective, the headroom analysis will clearly indicate what scope for growth exists in the current grid environment, and identify where investment is required to increase this.

Case study 2 – De-risking change

Sumerian was approached by a large international bank to investigate the performance and capacity of a post-trade analytics system, which comprises 3 grid partitions and over 2000 nodes. In particular, the IT team needed assistance to ascertain what the impact of proposed upgrade changes would be on performance and capacity; these included a switch to virtualisation, increasing the load by 100% and introducing a new server and database upgrade.

To achieve this, Sumerian captured infrastructure data from the application's underlying system components, baselining its current end-to-end performance and capacity and correlating it against business demand. By applying scenario modelling, Sumerian made significant findings that helped the team de-risk their decision making and plan the changes with confidence. For example, in modelling the impact of a switch to virtualisation, Sumerian's analytics concluded one server cluster was heavily imbalanced and through deploying virtualisation could reduce maximum utilisation by nearly 50%, giving them increased headroom and a much greater flexibility to meet changing business demands.

In terms of understanding the impact of growth by increasing the business load, Sumerian confirmed that there would be sufficient capacity to process the current volumes of portfolios, but during peak load periods, a bottleneck could appear in two key components. If the current load was to increase by a further 100%, the system would continue to function, but the time taken to process portfolios would increase by an order of magnitude at busiest periods. Finally, for the proposed server and database upgrade, Sumerian's modelling could accurately quantify that their introduction into the service would provide capacity and performance benefits by dramatically increasing the throughput of portfolios per hour to a far higher degree.

Furthermore, by correlating these baselines together the effects of business load on grid utilisation can reveal whether there are any direct correlations between business metrics (volume of trades and portfolios, type of analytics being requested, etc) and the effect this has on the grid's processing, compute time and performance behaviour. For example, analytics can extrapolate how CPU utilisation, I/O and latency are impacted by various business volumes and the composition of their mix. If a strong correlation is found, it produces a quantified relationship that can be used as a

predictive model to determine how both the grid and a business metric's behaviour will change under different variances. For example, if it is seen that the completion time for an overnight batch run is directly related to the type and volume of transactions contained within it, it may be possible to reconfigure grid partitions to improve the overall processing rate and thus shorten the batch's completion time.

Case study 3 – Relocating the grid

In work carried out by Sumerian, the trading arm of a large UK bank was planning to move their grid out of its current City of London offices to a more cost effective premises outside the M25. The core application infrastructure was to remain in London, with the grid engines being outsourced. The bank had worked with Sumerian on previous analytics projects, and had found its techniques to be highly advantageous in de-risking such large-scale change.

Sumerian captured packet traces from the existing network along with data from the grid's task reports, to build a model of the grid's current work loads. Sumerian was then able to forecast what the impact of additional latency would be between the outsourced grid engines and the core application infrastructure. The results of the analysis revealed that most applications would not experience any serious performance degradations with the additional latency; however, one application had a performance profile which placed architectural restrictions on the model to prevent data traversing multiple offsite locations.

Sumerian's analytics enabled the bank to plan which applications could be outsourced to the new grid location, and take appropriate action to keep the latency sensitive ones within shorter proximity. As a result, Sumerian enabled the bank to de-risk the project and ensure services were not impacted by the additional latency introduced.

Scenario modelling and forecasting

From understanding current capacity, performance and compute time, scenario modelling techniques can also be applied to gain precise answers for forecasting growth requirements or making predictions around the impact of architectural changes. The advantage of scenario modelling is that it can safely and accurately answer any number of "what if?" questions around the grid – essentially providing a "pre-sandbox" environment in which to test

out ideas before implementation. For example, in the case of predicting the effects of a 40% increase in a particular trade type, scenario modelling can predict the impact this will have on the grid's utilisation and capacity and whether there is adequate headroom to cope with the increased demand. In the case of infrastructure changes to the grid, scenario modelling can again provide quantified answers on whether investments will deliver against the desired improvements.





The path to optimisation

High performance grid computing is a critical element for today's investment banks. In order to achieve optimisation and maximise the potential of the grid's processing cycles, organisations should seek to build a strong dialogue and improve alignment with the business, to understand its priorities and gain visibility into the grid's current and future volumes. Without adequate understanding into where inefficiencies lie, organisations risk investing in areas where little or no business value will result. Instead, by using analytics' quantitative approach to baselining the grid and business environment to gain true visibility into capacity headroom and utilisation, organisations can begin to make steady, pragmatic improvements to achieve optimisation goals.

Analytics' ability to scenario model the grid provides even greater intelligence and responsiveness in which to fully understand not only the cost, but the risk and impact change initiatives will have on both the grid and trading performance. In applying an analytics approach to achieve this, investment banking IT teams can be confident in compiling a grid optimisation strategy that not only delivers against trading and cost objectives, but serves to increase the competitive standing of the bank and its future growth.

More information

For further information on Sumerian or to arrange a demonstration of our services, contact us on 0141 229 7580, e-mail us at info@sumerian.com or visit our Web site at www.sumerian.com

