

The Virtual Enterprise: Eliminating the Risk of Delivering Distributed IT Services

A Shunra Software White Paper



Executive Summary

The transition from development to production is a risky one. Distributed applications that look good on the developer's bench often under-perform in the production environment. Even apparently innocuous changes to an application can cause big problems when they're introduced into production. These snafus delay time-to-benefit, force costly eleventh-hour coding and infrastructure changes, and undermine IT's credibility.

This transition is so risky because the development environment does not adequately take into account the distributed nature of the enterprise environment. In the real world, services are delivered to geographically dispersed end-users across a complex enterprise environment. The specific nature of this environment – including the number, behavior and distribution of end-users – therefore has a direct impact on application performance.

Companies often attempt to ensure the performance of distributed applications in the real-world production environment by over-provisioning capacity or investing in various infrastructure enhancements. But these expenditures don't always prevent performance problems or produce worthwhile performance improvements. So precious IT dollars wind up wasted and performance problems remain.

A smarter approach is to pro-actively model the behavior of distributed applications in the real world throughout the project lifecycle. Applied from early development and QA through pre-deployment readiness testing and ongoing production management, this best practice ensures that target service levels are met consistently and efficiently.

The Shunra Virtual Enterprise uniquely supports this critical best practice. It provides a precise replica of the production environment – including remote offices, remote users and the delivery system that connects them. By applying the Shunra Virtual Enterprise throughout the project lifecycle, IT can find and resolve application performance problems before – not after – they are launched into production. They can better perform key lifecycle management tasks such as capacity planning, troubleshooting, impact analysis and the validation of proposed infrastructure investments. In other words, they can take the risk out of delivering distributed applications across the enterprise.

The Shunra Virtual Enterprise can be effectively applied to all types of distributed applications, including commercial ERP/CRM/SCM, eCommerce and custom development projects. It can support and rationalize initiatives such as data center consolidation, SAN/NAS implementations and disaster recovery. And it can ensure the successful introduction of new technologies such as VoIP, videoconferencing and Web Services into the enterprise.

As technological complexity continues to grow – and as the stakes of the game continue to rise – IT must eliminate risk from the delivery of distributed applications. No one should have to cross their fingers when they flip the switch on a new service or infrastructure modification. With the Shunra Virtual Enterprise, they don't have to.

The Forgotten Delivery System

The development of distributed applications is complex and challenging. It isn't easy to transform a long list of business requirements into working code. It takes a lot of skill and technology to ensure the quality of that code. Application functionality and usability must be repeatedly tested to ensure that users' needs are being met. And great care must be taken to manage code so it can be appropriately modified as business needs change.

Unfortunately, in the midst of all these development challenges, one critically important factor is often neglected: the distributed environment into which the application will ultimately be deployed.

This distributed environment is also very complex. It includes network infrastructure comprised of routers, switches, load balancers, packet shapers and network connections of varying capacity and distance. This network infrastructure connects multiple tiers of application servers, database servers, storage devices and other resources. These resources each have their own performance and configuration attributes.

Distributed applications depend upon this complex multi-tier environment as they execute their precisely choreographed sequences of requests, acknowledgements, queries and replies. Even a relatively subtle misstep or delay in communication between distributed components can substantially affect the end-user's experience with an application. The chain is only as strong as its weakest link. A weak flaw anywhere in the delivery system can cause the performance of the application to deteriorate or even fail.

That's why it's essential for IT to fully understand how its applications will behave in the real world of the distributed enterprise environment.

The Risk and Cost to the Business

Unfortunately, many IT organizations are not able to accurately model or predict the behavior of distributed applications in the real-world production environment. This inability exposes IT to a variety of significant business risks when deploying and managing such applications. These risks include:

Business interruption

When the performance of an application deteriorates – or when services fail outright – it can be very expensive for the business. Depending on the industry and the application, the cost of such interruptions can range from thousands to millions of dollars an hour.

Loss of credibility

When IT projects are delayed or fall short of user expectations, the credibility of the department and its management suffers. This loss of credibility isn't just a personal issue. It can have serious consequences ranging from lower budgets to an inability to get future initiatives green-lighted.

Delayed time-to-benefit

When problems in an application's ability to run effectively are discovered in production, it significantly delays the successful roll-out of that application. This postpones the anticipated benefits of the application – which can include new revenue, operational cost savings, improved customer service and/or mandated regulatory compliance. The cautious, incremental approach to production roll-outs that has often been used by IT organizations, that are unsure of how well their applications will perform in the real world, also tend to delay these benefits.

Unnecessary infrastructure investments

Because they know they're operating in the dark, many IT organizations try to give themselves some "wiggle room" by over-building capacity and/or investing in infrastructure solutions that promise to boost application performance. But over-building means over-spending – something few IT organizations can afford to do. According to Gartner, for example, over 60% of all companies over-spend on bandwidth by at least 20% because of this un-optimized approach. Worse yet, these expenditures often don't deliver hoped-for service improvements because the underlying problems have their root cause in the design of the application or in bottlenecks elsewhere in the environment.

Development over-runs

It's much easier and less expensive to fix problems in an application early in the development process. In fact, according to the National Institute of Standards and Technology (NIST), it costs 470 to 880 times more to fix a problem in an application at the end of development than it does early on. IT organizations that can only discover performance problems after an application has been deployed in the production environment will thus consistently over-spend on development.

Operations headaches

Applications that are not thoroughly validated prior to deployment on the production network can cause continuous headaches for network managers and other members of the operations team. These headaches include intermittent problems and “brownouts” with obscure causes such as low-speed PDAs accessing a system or undocumented data backups. The resulting fire-fighting and performance troubleshooting is costly and time-consuming. It also pulls staff experts away from more strategic tasks.

These risks and costs are increasing as businesses become more and more dependent on IT services for everything from supply-chain management to customer service. At the same time, the growing complexity of both applications and distributed environments is increasing the likelihood that unforeseen subtleties will adversely impact performance.

The bottom line is that IT must be able to model and analyze the behavior of applications in the real-world production environment throughout the development process. Otherwise, the business will remain exposed to the risks associated with the transition between development and production.


Best Practices for Eliminating Risk

Lifecycle best practices are essential for eliminating the risks associated with the delivery of distributed applications.


One important best practice is the creation of a baseline model of the distributed production environment. The various attributes of this environment – its infrastructure components, the number and geographic distribution of end-users, other existing application traffic, etc. – must be precisely modeled to reflect existing production conditions. This model provides a starting point for all future assessments of applications and infrastructure.


The next set of best practices applies this model throughout the development process. This should begin at the earliest stages of design, planning, and unit testing. In this way, any problems in the potential behavior of application code in the real world can be discovered and remedied early – before they become ingrained in the application’s design and therefore extremely expensive (or impossible) to fix.

Today:



Best Practices:



 Distributed application problems

When applied from the earliest stages of the application lifecycle, best practices performance modeling and analysis enables the discovery and remediation of problems before they emerge in the production environment – where they are extremely expensive to fix and create unacceptable business disruptions.

The model of the production environment should continue to be applied at each stage of development to ensure that potential real-world deployment problems are not allowed to creep into the code at any point. Thus, when the application is complete and has passed all of its use-case testing, it can also be put through a series of final pre-deployment tests to ensure that there will be no surprises when it is introduced into production.

Another set of best practices can be applied once applications are in production. If problems do arise in the production environment, modeling and analysis capabilities can be used to determine whether or not proposed remedies will actually bring about their desired effect – as well as to find out if they generate any unintended and adverse consequences. This is much more effective than experimenting with remedies in the actual production environment – where they can be very expensive to implement and where unforeseen consequences can result in costly and disruptive downtime.

Models should also be applied as required over time to validate and analyze any number of “what-if” scenarios associated with planned infrastructure modifications, projected growth in business demand, contingency planning and other IT due diligence.

Essentially, these best practices boil down to simple, good IT governance. As with any aspect of business (or of life in general), early detection is the key to avoiding and/or minimizing problems. Diligent modeling and analysis enable this essential early detection. They also give IT the ability to assess and validate technologies and initiatives before moving forward with them. So whether it’s a major new application roll-out or a supposedly “minor” modification in the network, IT must apply best risk management practices to protect itself and the business from unexpected consequences.

The Solutions Gap

Despite the obvious importance of taking the limitations and idiosyncrasies of the production environment into account during the development and testing of distributed applications, few IT organizations do so.

This has been due to a fundamental gap in IT's technology portfolio. This gap exists because none of IT's traditional tools bring the real-world characteristics and limitations of the enterprise production environment to bear on the design, development and testing of distributed applications before those applications are actually launched in that environment.

Developers work in local, dedicated environments that have none of the attributes of capacity, complexity, contention or distance that characterize the delivery system on which their applications will ultimately depend. They use load testing tools, but these don't reflect the complexity or topography of the distributed environment. They may also use some sort of simulation tools in the design of applications, but these typically do not replicate the actual production environment with any accuracy or granularity.

In other words, IT never knows exactly how distributed code will actually behave in the real world. So the deployment of distributed applications always involves unacceptable risk.

Once an application is launched into production, IT has a variety of performance management and optimization technologies at its disposal. But these fix problems after the fact. They don't enable IT to discover and fix problems before they impact the business. And if the root-cause of the problem turns out to be a flaw in application design, it's usually too late to do anything about it.

The situation can be likened to that of a pharmaceutical manufacturer. Developing an application is like developing a drug. You have to make sure the complex biological system into which you're introducing the drug is capable of handling it. You have to make sure that the drug can do what you want it to do and that it doesn't do anything you don't want it to do. You also have to know whether it will interact in any unforeseen ways with any other drug.

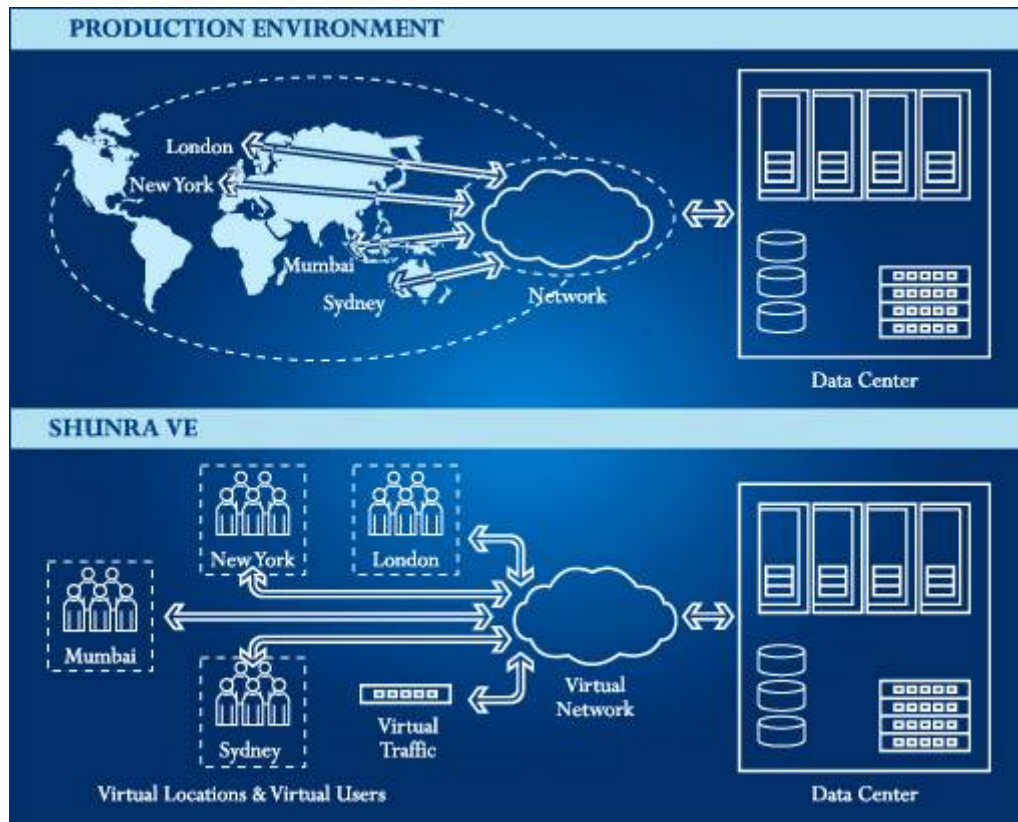
For IT, drugs represent applications and the human body represents the production environment. IT organizations simply don't have a way to discover whether or not their "drugs" will actually have their desired effect or not. They can't determine in advance if the "body" into which they are introducing the drug will be adversely affected by it. Nor can they discover any potentially adverse or fatal "drug interactions" before they go live.

In other words, the transition from development to production essentially requires IT to make a major leap of faith. That's not something we would accept from a pharmaceutical manufacturer, an automobile maker, or any other company introducing a new product into the marketplace. We expect them to have tested their products under real-world conditions before they release them. There's no reason for IT to be any different.

The Shunra Virtual Enterprise

So how exactly can IT address all these issues and risks? How can it implement best practices to take risk and cost out of the introduction of new and modified applications into the distributed environment? And how can it ensure quality and cost-efficiency throughout the rest of the application lifecycle?

The Shunra Virtual Enterprise offers a powerful answer. By fully replicating existing and/or projected conditions in the distributed production environment, the Shunra Virtual Enterprise enables IT organizations to adopt essential modeling and analysis best practices across the entire application lifecycle. It thus fills the solutions gap – enabling IT to aggressively take risk out of application development and delivery.



The Shunra Virtual Enterprise accurately replicates existing or projected conditions in the distributed production environment – including infrastructure, application traffic and the distribution of end-users.

The Shunra Virtual Enterprise is like a wind tunnel for IT. Just as an aviation company uses a wind tunnel to safely experiment with different aircraft designs under varying atmospheric conditions, the Shunra Virtual Enterprise empowers IT to safely experiment with different applications under varying production conditions.

IT departments can use this “wind tunnel” to test new or modified applications before deployment. They can use it after a problem has occurred to discover its root cause. They can use it to assess the impact of planned infrastructure changes or for discovering potential problems in any number of projected “what-if” scenarios. It can be used for business continuity planning, since it enables technology managers to determine in advance whether failover infrastructure is sufficient to support critical services. Its predictive capabilities can also be used to make smarter decisions about outsourcing, post-M&A infrastructure consolidation, and other high-impact business events.

The Shunra Virtual Enterprise is thus an indispensable resource for IT organizations dealing with the complexities and risks associated with the delivery of multiple business applications and services across ever-changing distributed environments. It enables IT to fulfill its service delivery mission with greater certainty, efficiency, consistency and speed. This makes it an essential component of the IT technology portfolio.

The Shunra Virtual Enterprise delivers all these capabilities by accurately replicating the current or projected attributes of the distributed environment, including:

Network infrastructure

The behavior of applications and services in the distributed environment depends to a large extent on the characteristics of the network. Networks have a certain amount of capacity at every point. They are also characterized by the geographic distances between their various infrastructure components, which introduce latency into the environment. These capacity limitations and latencies significantly impact application performance.

The Shunra Virtual Enterprise captures all these network characteristics and replicates them in its virtual model of the enterprise. It also allows network characteristics to be experimentally modified, so that technology managers can determine, for example, what will happen if a server is moved from one location to another or if a low-bandwidth link is put in place to support a new overseas office.

Application traffic loads

The performance of applications and services is also affected by the presence of other applications and service in the environment. Applications contend for the finite capacity available to them. So any effective model of the distributed environment must address these contention issues.

The Shunra Virtual Enterprise replicates existing application traffic patterns on the network. It also allows traffic patterns to be experimentally modified, so that technology managers can determine what will happen if email and web volume continue to increase or if a service like videoconferencing is introduced into the production environment.

End-users and remote offices

The performance of applications and services is also affected by the number and distribution of end-users in the environment. The behavior of end-users obviously has a great influence on application traffic loads. And it is the end-user's experience with the application that ultimately determines whether or not it is performing adequately.

The Shunra Virtual Enterprise replicates the number and distribution of end-users on the network. It also allows technology managers to experiment with changes in the end-user population. This allows them to take into account whether a new office will initially be staffed with 10 users or 100 users. It also enables them to determine what will happen if headcount in a specific office is increased by a certain amount. Just as important, it allows them to determine if any users anywhere on the network are actually experiencing unacceptable delays in application performance or outright application failure – which, after all, is the point of the whole exercise.

Benefits of the Shunra Virtual Enterprise

Because the Shunra Virtual Enterprise can be applied to so many of IT's most pressing challenges, it offers an extensive range of high-value benefits to the business. These benefits include:

Enhanced IT credibility

In business, it's critical to do what you say you're going to do. When applications fail unexpectedly, it undermines the credibility of both IT executives and IT as a whole. By enabling IT to more consistently deliver on its promises, the Shunra Virtual Enterprise boosts IT's credibility and ensures end-user satisfaction.

Improved operational efficiency

By enabling IT to discover and resolve problems in applications and infrastructure early and iteratively throughout the development process, the Shunra Virtual Enterprise eliminates costly re-engineering cycles, eleventh-hour re-coding, and disruptive troubleshooting "fire drills." This makes IT more efficient and allows technical staffs to focus on more strategic tasks.

Faster time-to-benefit

By pinpointing application design issues early in the project lifecycle, the Shunra Virtual Enterprise eliminates the delays that typically occur when unexpected problems are only discovered during deployment. It also eliminates the need to slowly roll out applications in multiple, incremental stages – as many IT organizations do because of the uncertainty of real-world performance. The Shunra Virtual Enterprise thus accelerates full production activity and allows the business to begin reaping the benefits of new applications and services more quickly.

Avoidance of unnecessary infrastructure investments

In order to ensure that critical applications perform at acceptable levels, many IT organizations provision excess capacity and/or invest in performance-boosting technologies. Unfortunately, these measures often can't compensate for basic flaws in the design of an application. With the Shunra Virtual Enterprise, IT can find out if additional capacity or performance-boosting technologies will actually help applications perform any better – or which combination of approaches is optimum – before it spends its money.

Predictable service-level assurance

A central mission of any IT organization is to deliver consistent, predictable service levels to the business. With the Shunra Virtual Enterprise, IT can fulfill this mission by pro-actively discovering and addressing issues that impact service levels before applications are put into production. And after applications are in production, the Shunra Virtual Enterprise enables IT to ensure that neither changing conditions in the production environment nor modifications in the application itself jeopardize target service levels in any way.

Prudent change management

Anyone who has been in IT long enough knows that the worst service failures often result from changes made to the environment by technical staff, rather than from the failure of some hardware device. That's because such changes often have very unexpected consequences. With the Shunra Virtual Enterprise, on the other hand, planned changes can be performed in a virtual environment first. This allows any unexpected consequences to be safely discovered before they occur in production, so that they can be properly addressed and avoided.

More precise capacity planning

Many companies overbuild their networks because they're just not certain how future growth (in business activity, in staff, in the number and location of facilities, in the number of applications) will affect bandwidth needs. With the Shunra Virtual Enterprise, capacity planners can execute a wide range of "what if" scenarios to assess exactly how projected growth will impact the network. By modeling these scenarios, IT can therefore avoid purchasing too much or too little capacity for anticipated needs.

Faster troubleshooting of intermittent problems and other network issues

Intermittent application performance problems present a particularly tough challenge for IT engineers. It's hard to troubleshoot an issue that you can't see, so eliminating such problems often takes weeks or months as engineers keep trying to examine it during each sporadic occurrence. In these situations, the Shunra Virtual Enterprise provides an extremely valuable controlled environment where such problems can be frozen in a sort of forensic "snapshot"—so that engineers can fully analyze them and formulate effective fixes. The Shunra Virtual Enterprise also affords them the luxury of seeing if their proposed fix actually solves the problem or not before they implement it in the production environment. This same principle can be applied to other types of network issues, since it's often easier to find a problem than it is to figure out exactly how to fix it.

These benefits are extraordinarily important for IT as it finds itself called upon to deliver and support a growing number of revenue-critical services in an increasingly dynamic distributed environment. The reliability, cost-efficiency and speed that the Shunra Virtual Enterprise provides are essential for any IT organization hoping to successfully meet the evolving needs of the business within today's tight resource constraints.

Success in the Face of Complexity

There are plenty of reasons for IT executives to lose sleep. The business continues to turn to IT to help it better market its products and create more closely bonded relationships with customers. Technology has become the crucial element for squeezing cost and time out of the supply chain and for ensuring compliance with increasingly stringent regulations. And at the same time as IT is being pressed to design new business solutions, it must continue to maintain the health of existing ones. With multiple generations of technology to manage and new generations on the horizon, IT has become more stressful than ever.

But IT executives shouldn't have to worry about whether the application they've just spent nine months and ten million dollars building will actually work when it's launched in production. They shouldn't have to suffer grief over a critical business system that keeps freezing up at the end of the month. And they shouldn't have to make decisions about major network purchases based on pure instinct.

With the Shunra Virtual Enterprise, they don't have to. The Shunra Virtual Enterprise provides IT executives, managers and technical staff with the safe, controlled environment they've always needed to test, analyze, project and predict the behavior of distributed applications and infrastructure in an intelligent, prudent manner. By incorporating the Shunra Virtual Enterprise and its associated best practices into every aspect of planning, development, deployment and management, IT can eliminate significant sources of risk and cost while accelerating time-to-benefit and rigorously protecting its credibility.

About Shunra

Shunra provides solutions that empower organizations to address service level and performance concerns up front – before deployment. The Shunra Virtual Enterprise (Shunra VE) solution provides accurate, highly granular insight into how networked applications will function, perform and scale for remote end-users. It creates an exact replica of the production network environment, allowing IT professionals to safely develop, test and experiment with applications and infrastructure before rollout, and effectively plan for growth and change. With solutions tailored for networking and performance professionals, software developers, and quality assurance staff, Shunra VE facilitates collaboration across all IT disciplines – so IT organizations can quickly and more efficiently uncover and resolve problems before they impact the business. This results in more timely, higher quality and cost-efficient IT services, and the ability to “Deliver IT with Confidence”

Solutions for Any Enterprise

More than 1500 customers, including hundreds of *Fortune* 1000 and Global *Forbes* 2000 organizations, from financial institutions to manufacturing companies, retail, energy, media companies, as well as independent hardware and software vendors and telecommunications service providers, have gained measurable returns from Shunra’s solutions. Among them are: 3M, Boeing, Cisco, Dow Chemical, EMC, FedEx, General Electric, General Motors, JPMorgan Chase, Kelly Services, Merrill Lynch, Motorola, Nestlé, Pitney Bowes, and Vodafone.

Corporate Information

Shunra’s headquarters are located in New York City and Kfar Saba, Israel, with worldwide offices in Singapore, UK, The Netherlands and India. Shunra is also supported through a global network of channel partners.