Intellyx White Paper

Bring Web Scale to Enterprise IT Operations Management



Customer-driven digital transformation, the explosion of big data, and the rise of enterprise web scale technologies are all raising the bar on the challenges that IT operations management faces every day.

Business stakeholders from the C-suite on down now realize that real-time performance of their software-driven enterprise is essential for their profitability. Whether it be seamless mobile banking, delightful ecommerce experiences, or blisteringly fast in-game play, web scale performance is now the coin of the realm.

IT operations management must rise to this challenge – leveraging real-time, web scale tools and technologies to resolve issues and guarantee the performance so critical for the success of today's modern digital enterprise.

The World of Web Scale

A mere handful of years ago, the world of enterprise IT consisted of legacy systems, physical servers, and inflexible architectures. Cost constraints ruled the day – do more with less, keep the lights on, and struggle for relevance in an era when IT didn't matter.

Then along came Google. And Amazon. And eBay. Eventually Facebook, Netflix, and scores of other technology-based companies joined in, rethinking what IT was all about. Many of these twenty-first century darlings drove innovation in hardware, networks, software, and the architectures that made everything work.

Welcome to the world of web scale.



In spite of these astounding advances in technology capability, the denizens of enterprise IT continue to struggle with legacy systems, legacy software, legacy...everything. Their world still centers on constrained resources – not enough speed, not enough scale, not enough money to go around. Keep the lights on and maybe – just maybe – you'll get to do a bit of innovating as well.

The web scale world is just the opposite. These companies had cracked the nut – speed and scale in abundance, with no ball-and-chain legacy in sight. Not coincidentally, their pace of innovation has reached unprecedented levels, as web scale market caps surpass those of traditional enterprises, many of whom spent decades to grow to their current size.

Web scale, therefore, represents both the technology itself, as well as a way of thinking about how to manage and grow an innovative business at scale.

Enterprise executives around the world are now thinking to themselves: we need that. Now we know it works. How do we bring the world of web scale to the enterprise?

Cloud computing, of course, is part of the answer – but today, the story goes well beyond the cloud. The challenge across enterprises today is *digital*.

Digital represents an accelerating shift in customer preferences and behavior, which are now driving enterprise IT decisions like never before. As a result, digital transformation initiatives are driving enterprises to operate at web scale.

Digital also raises the bar on customer expectations of performance and security, as the number of touchpoints, applications, and environments explode. Somehow, however, IT shops must manage this rapidly changing technology landscape. Enterprises must bring the web scale context to IT operations in order to remain competitive.

Because today's digital applications are increasingly distributed and ephemeral, running both on premise and in the cloud, IT ops will never be the same again.

Bringing Web Scale Thinking to Enterprise IT Ops

DIGITAL RAISES THE BAR ON CUSTOMER EXPECTATIONS OF PERFORMANCE AND SECURITY, AS THE NUMBER OF TOUCHPOINTS, APPLICATIONS, AND ENVIRONMENTS EXPLODE.

As enterprises progress with their digital transformation initiatives, they typically implement elements of web scale within their IT organization. However, you can't take a century-old bank and expect to turn it into Google – nor would you want to.

Fundamentally, web scale means something different to a traditional enterprise than to the web scale companies themselves. We call this perspective *enterprise web scale*.



Web scale companies, of course, had the luxury of being able to architect their environment for web scale from the beginning. Enterprises, however, are not so lucky. Such companies scale by adding new capabilities onto existing ones, as retiring legacy assets is rarely practical or cost-effective.

Becoming web scale for an enterprise, therefore, typically means adding web-scale capabilities to existing on-premise assets – a ticket to the world of web scale to be sure, but also the recipe for additional silos – both technological as well as organizational.

Perhaps an enterprise leverages a public or private cloud in addition to existing on-premise capabilities. Another common pattern is for a company to build out big data capabilities (either in-house or in the cloud) that must now integrate both with the sources of data as well as with the business intelligence, visualization, or other analytics tools that consume and display the results of the big data processing.

In spite of the silos so typical of any large enterprise, the digital priorities of the organization are *end-to-end*. All elements of a customer-facing app, from the user interface all the way to the underlying systems of record, must perform all the time, and at top speed.

In fact, the end-to-end requirement – both technological and organizational – is perhaps the most transformative aspect of digital transformation for many organizations.

This combination of technology silos and customer-driven performance requirements upends traditional IT operations management. Existing management tools align with technology silos – and adding web-scale capabilities only exacerbates this problem.

And yet, the enterprise context for web scale requires cross-cutting visibility as well as root cause analysis of issues that may involve different elements of the end-to-end digital transaction.

Managing Web Scale: A Big Data Challenge

Traditional IT ops management tools simply aren't up to the enterprise web scale challenge. The reason? *Monitoring enterprise web scale environments is itself a web scale problem*. Even for enterprises with limited web scale initiatives, the end-to-end monitoring requirement nevertheless calls for both web-scale thinking and web-scale architectures.

MONITORING ENTERPRISE WEB SCALE ENVIRONMENTS IS ITSELF A WEB SCALE PROBLEM.

As a result, managing IT ops is becoming a big data challenge – but just how big are big data anyway? There is no official definition of big data, but one common definition is *data sets that are too large for traditional tools to store, process, or analyze.*

Big data, therefore, are a moving target – as tools mature, the threshold for big data continues to grow, remaining at the boundaries of our ability to deal with increasingly large data sets.

Web scale environments are always pushing this threshold, of course. Managing enterprise web scale environments also continues to exceed the capabilities of traditional enterprise tooling – and will continue to do so for the foreseeable future.



In fact, generating several terabytes of ops data a day is increasingly common. A terabyte is difficult to envision – one trillion bytes, corresponding to the text in a billion-page book, perhaps. A few years ago the Library of Congress famously contained 20 terabytes of information – although that number is certainly higher now. Then again, terabyte hard drives run around \$50 these days, showing just how common such quantities of information have become.



The typical enterprise generates more ops data in one day than in the entire Library of Congress

Root Cause Analysis: The Core IT Ops Management Challenge

Perhaps the most important big data challenge for enterprise web scale IT ops management is *root cause analysis*. Enterprise web scale requires that IT ops monitoring be real-time, but root cause analyses must look at past behavior as well in order to uncover underlying causes of anomalous events.

As a result, IT ops tooling must deal with both real-time behavior as well as increasingly large and complex archives of event data – a big data challenge in and of itself.

Furthermore, as with so many modern big data problems, such data come in different types with different levels of structure. Some of the raw data are unstructured log files, while in other cases, the information the IT ops tool requires is the result of structured queries or various types of API calls.



Root cause analyses that involve many terabytes of data present unfamiliar challenges to the IT ops organization. It's important to merge real-time analysis and analysis of historic data – even seconds after the fact – for effective root-cause analysis.

In today's digital age, real-time processing is essential for nimble, responsive business – but it's never perfect. Uncaught errors cascade, and the faster the processing, the more errors crop up.

Furthermore, most big data analysis seeks to identify correlations among data. To determine causes of anomalies, however, data correlations aren't good enough. We must separate causes from symptoms over time in order to say when systems are behaving differently and why.

IT ops must also correlate data across silos and diverse infrastructure components – a particularly difficult task for modern IT monitoring tools, and entirely impossible for traditional enterprise tooling at the velocity and quantity of data that enterprise web scale generates every minute of every day. WE MUST SEPARATE CAUSES FROM SYMPTOMS OVER TIME IN ORDER TO SAY WHEN SYSTEMS ARE BEHAVING DIFFERENTLY AND WHY.

Traditional tools simply aren't able to correlate data across different monitoring silos, different structures, or different time periods, and with the dozens of monitoring tools in a typical enterprise, simply integrating this mishmash of tooling won't address such challenges.

Furthermore, anomaly detection does not consist entirely of identifying inadvertent problems. Hackers cause anomalies as well – and not only is it critical to catch attacks immediately, but any IT ops monitoring tool must be able to uncover the tracks of hackers who are trying to hide those tracks.

Finally, simply applying patches to vulnerabilities is necessary but not sufficient, as there is always the possibility a patch is ineffective. Don't assume you've patched every vulnerability. Ops teams must continue to verify that previously resolved threat vectors remain resolved on an ongoing basis – yet another big data challenge.

Bringing Web Scale Technologies to Bear

If the enterprise retains its historic constraints on its ability to store and process large quantities of data, deep analyses of heterogeneous performance data may not be possible – especially when the IT organization already runs numerous, disparate monitoring tools. Instead, it's essential to leverage the web-scale context of virtually unlimited data storage and processing capability in order to be successful with IT ops management in an enterprise web scale environment.

Such storage and processing comes at a cost, to be sure – but the financial context must also move to the pay-as-you-go, pay-for-what-you-use model of cloud computing. Licensing IT ops tools by the core or by the user is no longer an appropriate financial model now that enterprises are participating in web-scale computing.



Furthermore, enterprises must leverage the experience and hard work of web scale companies in order to succeed with their own web scale efforts. Fortunately, the open source community provides an entrée to any enterprise looking to succeed with web scale technologies.

<u>Hadoop</u> is one of the best known open source efforts to bring web scale to enterprise IT shops, but it is merely a harbinger of greater technologies on the way. Today we have a plethora of open source initiatives that one way or another distill the best practices of web scale.

IT ops management vendor <u>Rocana</u> is riding this wave. They built their web scale ops management technology by leveraging the Hadoop File System, <u>Kafka</u> message broker, <u>Solr</u> enterprise search server, and <u>Impala</u> analytic database for Hadoop – all open source projects from the Apache Foundation.

Other vendors can deal with big data ops management challenges, but Rocana stands alone as having a platform built from the ground up for web scale ops management challenges. ROCANA STANDS ALONE AS HAVING A PLATFORM BUILT FROM THE GROUND UP FOR WEB SCALE OPS MANAGEMENT CHALLENGES.

As enterprises increasingly implement web scale technologies – and for any company, enterprise or not, who is web scale today – leveraging a fully web scale IT ops management platform like Rocana's is becoming increasingly important. In the future, all enterprises will have to approach IT ops this way.

Rocana is an <u>Intellyx</u> client. Intellyx retains full editorial control over the content of this white paper. Image credit: <u>Carol M. Highsmith</u> (CC BY-SA 3.0).