

[WP-022]

## Whitepaper: Healthcare IT Grid?

Risks or Issues Surrounding the Development and Maintenance of Healthcare Grids

**Version 050817**

August 2005

Author: Bob Radvanovsky, [rsradvan@unixworks.com](mailto:rsradvan@unixworks.com)

*(A special thanks goes to those listed for being my “sounding board” on this project.)*

Copyright © 2005 Bob Radvanovsky. All rights reserved.



## Limited Liability Statement

---

In no event shall the author(s) be liable for any found errors contained herein or for any direct, indirect, special, incidental or consequential damages (including lost profit or lost data) whether based on warranty, contract, tort, or any other legal theory in connection with the furnishing, performance, or use of this material.

The information contained within this document may be subject to change without any notice. No trademark, copyright, or patent licenses are expressly or implicitly granted (herein) with this whitepaper.

Documentation pertaining to any security-related technical or proprietary information, its data and all information provided and contained within this document is considered proprietary in nature and subject to copyright protection and is intended solely for use by its owner. Additionally, this documentation is solely for discussing a possible, and/or proposed critical infrastructure protection (CIP) security issues, and is not dependent upon any specified infrastructure, architectural condition or its issue(s). Source information is through observation from previous circumstances, related online documentation, and discussions with colleagues within the security field.

All other product names mentioned herein are trademarks or registered trademarks of their respective owners. NOTE: Any names not outlined or mentioned above are fictional in nature; as such, any relation to any name or trademark (if any) is purely coincidental.

## Introduction

---

The slow introduction of “*high performance computing*” (also known as “*grid computing*”) has been in use by high energy physics laboratories and federally-controlled National laboratories since the late-1990’s. Grid computing isn’t (per se) a new topic, but rather a modification of an already existing one: massively parallel processing that is split, demographically, based on specific functions of the overall system (and its process).

Strangely enough, Corporate America has not embraced grid computing as some would have hoped. Indeed, this newer technology, esp. with several of the critical infrastructure sectors, namely, transportation, finance and healthcare – all appear to be hesitant in their embracement of this newer technology. One stipulation is that these sectors have been rumored/speculated to be the most resistive to adaptation through change, due to (primarily) their inherited nature of their sector’s primary business function and role: if they “go down”, either people die, and/or companies loose a lot of money. It’s that simple. Companies don’t like loosing money, and they aren’t in the business of charity (can’t remember who told me this, but it was stated several years ago from a former work colleague from a larger healthcare provider in the Northern region of the United States).



Be that as it may, these companies have been very resistive to change in just about every function, partially due to union-related or (unions abhor change, too, esp. if it means loosing control over that company or sector, or if jobs are at stake – a prime example is the divergence of two (2) major, significant unions splitting from the AFL-CIO) or labor-related reasons, or other reasons such as inclusion into an already tightly infused and overtly ingrained and very complex infrastructure. Many critical infrastructures are considered critical due to their complexities and amount of funds required in keeping them up-to-date. Same goes for computing power, too.

However, despite whatever excuse that they may have (and rightfully and costly just so), does not mean that they should not consider alternatives. One of these alternatives is grid computing. It offers fault tolerant environments such that, if Texas were to go \*PIFF\* due to a hurricane, or just simply vanish, the computing environment would continue to operate as fault tolerance is a safety feature of grid computing. There are other reasons or factors for wanting, or suggesting, utilizing grid computing, some cost-related, others to slowdown or mitigate the effects resulting from offshoring of IT workers and personnel.

One such major or significant reason for introducing grid computing is cost – but from several perspectives:

- reduced cost of operations from a fault-tolerant, centralized data center;
- improved efficiencies based on regional usage and overall cost;
- improved efficiencies of an established standardized data model;
- improved availability of data for auditing and tracking capability purposes;
- improved tracking capabilities of suspected individuals by law enforcement;
- redundancy and centralization of data; and,
- accessible through secured communications channels

Personally, the other remaining critical infrastructure sectors could also benefit from this technology from a much larger wide-scaled implementation, but since the topic highlights the healthcare sector, and since it has been riddled with a plethora of political geometries that have embroiled the industry for over 10 years now, this would only appear fitting to introduce the concept of grid computing to this sector first.

## **Reduced cost of operations from centralized data center**

---

Probably the largest benefit of grid computing is its ability to appear as a centralized source, but in reality, is load/data-balanced across a distributed, geographically decentralized area (say, the ENTIRE United States), and break it into 5 or 6 regional areas. If you were a healthcare provider from Chicago, that required information from one source, you would simply access data through your regional provider that serviced the Chicago metropolitan areas. If the regional data center that serviced Chicago were to go down due to a tornado, flood, or even a security breach, would be distributed to other regional data centers based on load at that time, and location (if the healthcare provider was closer to an adjacent regional data center to the East, rather to the one to the West, it would choose that, depending on load and server availability).



So then the question remains, why does each and every healthcare provider have its OWN data center facility? This seems a tad bit inefficient of resources, and most times, their “technical staff” consist nothing more than medical staff who have been trained, or are self-trained for 1 or 2 specific applications used by that healthcare provider. How then, can CIOs from these healthcare providers state that they have qualified and trained personnel to handle their healthcare personnel’s needs? It simply comes down to the fact that these people are trainers more so than support personnel, as most times, they usually perform the “*one armed lift*” (this means picking up a handset receiver and calling, via telephone, the support number of the actual software or hardware provider or vendor); how is this supporting your staff?

Since healthcare providers want to reduce costs, this would be one of the most significant methods of cost reduction: removing staff who are of little benefit to the company, and who offer nothing more than some training experiences to other staff members.

The regional centralization would allow healthcare providers of all sizes to focus on their primary business function and role: servicing patients. If I fail to understand or comprehend why a healthcare provider keeps or retains personnel and staff members who offer very little in terms of cost-to-benefits ratios (overall) to the company, why then do they remain employed with that healthcare provider if they do – *nothing*? Perhaps I am off with this statement, but too many healthcare providers have too much “dead wood” that specializes in ONE function or application of ONE system. Most healthcare providers today have hundreds, if not thousands, or varied applications servicing differing needs, functions and specialties of the healthcare industry. So what’s wrong with this picture?

A streamlined, consolidated set of data centers for 5 or 6 regions, with a centralized data repository location for ALL DATA – is the preferred method. EVERYONE has access to EVERYTHING – period.

## **Improved efficiencies based on regional usage and overall cost**

Another advantage of utilizing a healthcare grid might be that Chicago bogs down the Midwest Regional location, and that future connection requests would be automatically redirected to adjacent regional data centers to the East, West, North and South. For regional data centers that are at opposite extremes of the country, could have cross-ties to other, possible lesser populated regions, which could accommodate for expansion or overflow processing from those that might need those resources.

This isn’t an unusual concept, and has been around for almost 10 years now through networked load balancing. Essentially, this principle takes multiple servers (web-based, application-based, terminal-based, etc.) and balances based on network load usage, network availability, security, natural disasters, connections failures, etc. All of this is controlled – regionally – by several data center hosting providers that provide network, hardware and/or software co-location services, or networking redirection capabilities off Tier 0 Internet connection pipelines.



## **Improved efficiencies of established, standardized data models**

This subtopic confuses me the most, yet raises a ton of questions in the process. It can be quietly simply be put and summarized into ONE question: how many times or methods does a field or form contain a “last name” of a patient, or can be displayed? Meaning – just how many different methods does an individual’s last name need to be displayed, formatted, “*enriched*” (I loved hearing that term “*data enrichment*” – have absolutely no idea what it does, other than I can surmise it means manipulation of the data into yet another form that is displayed and is charged for – huh?) for final output in UHC-1234-01 (this form number doesn’t exist, and was made up to show the stupidity of the over burgeoning amount of forms and paperwork that must be filled out today) form.

Let’s put it in another way: if every state were to come together and agree on ONE form that patient registration were to have filled out, this would eliminate much of the hubbub about if we had filled out the correct forms or not. For instances, forms from the State of Illinois might not be compatible with forms from the State of Nevada or the State of California. Why? They want to control their own data. It’s pretty simple. If data were formatted in their approved and adopted data model format, they can lay claim to that data as being “theirs”. If the data model were to be standardized into ONE data model, taking into account each and every nuance and minisculistic and idiotic piece of data that were to every be devised – wouldn’t that make life a tad bit more easier for everyone?

## **Improved availability of data for auditing/tracking purposes**

Because of HIPAA (in the healthcare industry, they only have 5-lettered swear words, and this is the mother of them), the federal government has instilled (if not demanded) the need for auditing capabilities of their patients, how much they have paid towards their final bill, what they owe, etc. etc. Of which, healthcare providers “*harvest*” (another marketing buzzword that I love which can be summed up colorfully to explain the term “*data mining*”, which tends to sound more harsh and roughened) to keep demographical information about current and past patients, visitors, spending habits, and what they should target for future patients, as well as those who have visited them. This makes perfect sense – target market your audience, right? OK, if you have to go to 12 different locations to get ALL of that data would make life a tad bit more complicated, wouldn’t you agree?

If data were centrally located for all healthcare providers, thus would make the sharing of healthcare information easier, esp. at the flick of a switch, or click of a web-based mouse button, can move data, or allow the viewing of data to be more systematically controlled based on levels of control, levels of access, etc. Essentially, a doctor would be allowed or permitted to view only his patients, and not that of fellow constituents; conversely, visa versa with his colleague’s patients.



However, with a centralized data repository, if Doctor A wanted Doctor B to have viewing privileges of Doctor A's patient, the doctor would send an automated email to Doctor B requesting acknowledgement, along with any legal statements to view, levels of understanding, etc. Once done, Doctor B could view Doctor A's patient's information *without* having to (per se) request the information. It could come from the same source – the healthcare grid.

## **Improved tracking capabilities of suspects by law enforcement**

OK – this may sound too much “Big Brother”-ish, but with the passing of certain law enforcement provisions within the USA PATRIOT Act of 2001, this is exactly what law enforcement wants – and needs. Secondly, this isn't necessarily something that would invade upon the privacy of individuals. If, for instance, you had a habitual drug addict, you regularly stole drugs, money, and other possessions, law enforcement could keep track of that individual's habits while under the care of the healthcare provider. This would allow for easier tracking capabilities, again, from a centralized data source.

The data stays “locked” until a court order releases, or makes it available for viewing by law enforcement personnel. Yes, this would make it easier for law enforcement to keep tabs on people, but would allow for great efficiencies of the system as a whole. One negative side-effect would be that the federal and state governments might utilize such a tool to dictate how much healthcare we would be authorized to have, and would allow for greater favoritism, quite similar to what some have stated is occurring elsewhere throughout the World, outside of the United States.

## **Redundancy and centralization of data**

I may have elaborated this concept more fully in the previous topic, but essentially, a centralized data source would allow for reduced paperwork, reduced forms to be filled out, standardization of data entered and utilized, and would allow greater sharing of data (properly and correctly) based on proper validation methodologies used by our military. Essentially, the “*Need to Know*” classification would be introduced, as the data/information would be considered an unclassified form of data/information, but would be sensitive enough to justify the necessary precautions of safeguarding the privacy and data retention capabilities of individual's identity.

Secondly, since we're talking about identity, having a centralized data source would drastically reduce the likelihood of the risk/threat of loss of one's identity through identity theft. Obviously, there would be a needed greater control of the availability of the data, viewing rights and privileges of that data, etc. All of this would take into consideration who has the “need to know” what about someone.



## **Accessible through secured communications channels**

---

The noticeable trend that I am seeing is just that. We are becoming more and more geographically distributed throughout an area, region, and our country. Thus, the need to find a much secured method of communications would drive costs down, but would still satisfy the increasing need for always-online communications. We have the Internet. OK – now what? How do we do that over an unsecured network?

Simple: trusted computing. Intel in early August, 2005 released information about the Trusted Computing Module (TPM), that would interface with the motherboard and existing processors, thus allowing for a secured connection over an unsecured network, without the need of requiring an individual of remembering methods of authentication. Passwords will never go away, but the method by which someone was to connect to a secured environment – might change. TPM through the use of Trusted Computing architectures, combined with other operating system and application environments, would allow for a grid-like environment to exist.

Now comes the next question – would one utilize their PC? Perhaps, or perhaps not. I foresee the use of thin clients making a HUGE comeback; esp. with the risks associated with people being able or is capable of storing off patient records, financial data, etc. Therefore, *smart* thin clients – are destined to come back. It is inevitable.

## **Conclusion**

---

So where is this all going? Probably within the next 10-20 years, several key industries will begin to see – and profit – from the utilization of this emerging technology. Grid computing is an inevitable feature that will allow greater flexibility, greater availability, greater security, and greater capabilities (as a whole) for all involved.

