ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEAPP02

MODERNIZATION CHALLENGES FOR THE IT INFRASTRUCTURE AT THE NATIONAL IGNITION FACILITY*

A. Casey, P. Adams, M. Christensen, E. Ghere, N. Spafford, M. Srirangapatanam, K. Tribbey, K. White, R. Vadlamani, D. Yee, Lawrence Livermore National Laboratory, USA

Abstract

As the National Ignition Facility (NIF) enters its second decade of full-scale operations, the demands on all aspects of the Information Technology (IT) infrastructure are becoming more varied, complex and critical. Cyber security is an increasing focus area for the NIF&PS IT team with the goal of securing the data center whilst providing the flexibility for developers to continue to access the sensitive areas of the controls system and the production tools. This must be done whilst supporting the interoperability of controls system elements executing on legacy bare metal hardware in an increasingly homogenized virtual environment in addition to responding to the user's requests for ever-increasing storage needs and the introduction of cloud services. While addressing these evolutionary changes, the impact to continuous 24/7 Shot Operations must also be minimized. The challenges, strategies and implementation approaches being undertaken by the NIF&PS IT team at the NIF to address the issues of infrastructure modernization will be presented.

DEFINING THE PROBLEM

With NIF routinely executing over 400 hundred experiments annually, there seems to be little issue with the IT infrastructure used to run NIF. However, ever since 2009 when NIF transitioned to a user facility, technologies have been changing at ever faster rates and the infrastructure has not been able to keep pace with them.

What was state of the art in 2009 is often considered obsolete in 2019. For example, in 2009 servers based on the Solaris 10 OS comprised most of the compute power in the Data Center but today it is difficult to find System Admins with the skillset to administer those hosts.

Secondly, NIF is first and foremost a research facility and maximizing operational availability is key to meeting the shot rate goals. However, this does conflict with the ability of the IT team to make the necessary evolutionary changes to the infrastructure in order to keep the facility operating as securely and as efficiently as possible.

REVISIT THE DESIGN

Before starting the process of updating the infrastructure, the decision was made to validate the current IT infrastructure against a set of standards that could be used to benchmark the current state and thus form the basis of an implementation gap analysis. As cyber security has become a critical part of IT operations and the risks highlighted by publicized incidents such as the Stuxnet worm and the City of Baltimore ransomware attack, the Center for Internet Security (CIS) Top 20 Controls and Resources [1] were chosen to provide the security framework.

IT Data Management

The basic CIS controls stipulate the need to actively manage all of the enterprise's IT assets through inventory control and configuration management. By analyzing the data produced by the IT tools monitoring aspects of the system, such as network traffic for example, making investments in new tools where data was not available, and then aggregating it all in a vizualizations tool, it was possible to develop a view into the infrstructure that could be then used to guide the next steps in the modernization process.

Industry Standards

Over the last 10 years, the NIF has evolved as operational needs have changed and as technologies have been introduced to improve performance and to make it easier for users – both internal and external – to do their work on a day-to-day basis.

Verifying that these changes do not break any of the standards utilized when the original infrastructure design was laid out, and that they are still applicable ten years later, was the next step in the process.

As the NIF's IT infrastructure encompasses Industrial Controls Systems (ICS) is part of the supervisory control and data acquisition (SCADA) network, a critical part of this effort was to ensure that the network design still conforms to the Purdue Enterprise Reference Architecture (Fig. 1). The goal has been to ensure that the paths from the institution network (level 5) to the controls themselves (level 3 down to level 0) are understood and its configuration managed. This addresses the concern that over time, "configuration drift" has added connection paths that are not desired.

Aside verifying design standards, effort was spent on ensuring that hardware and software was installed and configured as close to manufacturer recommended configurations as could be achieved. For example, a significant number of servers are based on the CISCO Unified Computing System (UCS) [2] platform and the installation was done in accordance with Cisco Validated Design (CVD) to ensure maximum compatibility with operating system, database and storage components.

Risk Assessments

The final part of the prcoess was to reflect on what had been learned previously. A lot of information was gathered in validating the infrastructure design and that data was then used to create a risk model for each of the CIS

* This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, Release # LLNL-CONF-790677

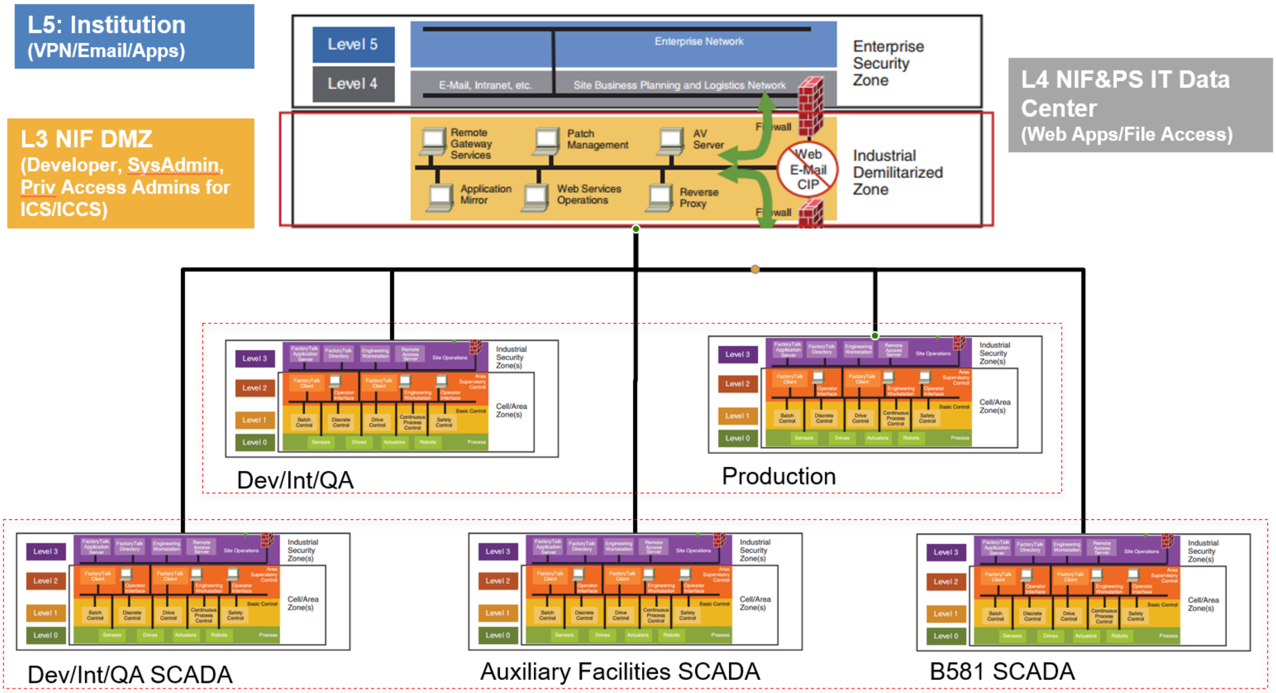

Figure 1: One of the considered options for NIF Cyber Security Network Model is to separate each of the environments below the DMZ according to the Purdue Enterprise Reference Architecture standard.

controls. The assessment was based on the design, the implementation of the control, the likelihood of the control being circumvented, and the impact to the facility if the control failed. This objective look at the infrastructure then gave an analytical way to talk to all stakeholders about risk without resorting to hyperbole, helped prioritize where future investment should be made and just as importantly, provided a benchmark against which improvement, and hence return on investment, could be measured.

THE PRODUCT LIFECYCLE

Development of software has changed significantly over the years. There is a lot more focus on the development lifecycle and how teams roll out new capabilities to users while minimizing the disruption that this causes.

The provision of IT services has not always followed such a model.

In order to introduce new capabilities, most IT teams will follow a standard process of requirements gathering, designing a solution, and reviewing it. However, the review should not be limited to a technical assessment of the solution. It also needs to include cost / benefit analysis of the update. Trying to keep up with technology is a costly venture and there needs to be a tangible benefit to the customers or the IT team before committing resources to it. The benefit can be risk reduction, operational efficiency or less scheduled down time but it needs to measurable.

The next phase of the lifecycle is to work with the controls and applications developers to roll the update through the development environments to validate the impact of the update. The motivation for this is to minimize the risk of the update having a negative impact on Production. Very few facilities have the luxury of replicating production in an offline environment but regression testing and monitoring builds confidence that the update does not have any unexpected side effects.

The final phase is to continually monitor the infrastructure. Using the SW development life cycle as a template, the monitoring covers the components of the infrastructure such as storage, network, and compute (analogous to unit testing in SW), the applications themselves (integration testing) and then performance monitoring. These monitoring "views" into the system give the IT team the ability to baseline behavior, validates that an update does what was expected and allows the team to see off normal behavior before the users do with the goal of resolving issues before they impact operations.

A high-level representation of this process flow is shown in Fig. 2.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WEAPP02

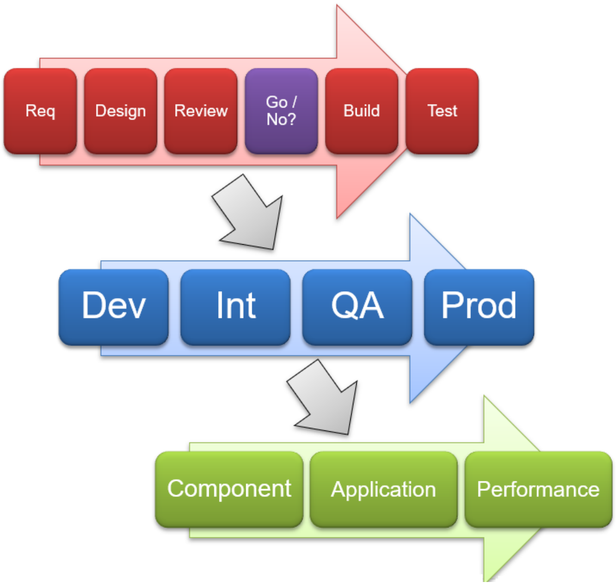

Figure 2: The IT product development lifecycle with a focus on ensuring realizing a return on investment.

COMPONENT HOMOGENIZATION

When the NIF&PS IT infrastructure was initially installed, the system was designed as hundreds of networked bare metal servers. These servers required a highly skilled team of Sys Admins to monitor and maintain them. As capabilities were added, so was new hardware which was not necessarily the same make and model of what was installed previously. This further complicates the maintenance of the system as it introduces; the need for more spares to cover all of the hardware models, management of different patching regimes to deal with varying levels of operating systems and other installed software, increases the load on the Sys Admins as they have more touch points (even with the assistance of scripting tools) and as each environment is different, lowers the Sys Admins agility and productivity.

NIF&PS IT followed the path of homogenization to address these short comings. Utilizing the CISCO UCS platform, a significant multiyear effort was begun to transition from the bare metal architecture to a component based virtual architecture. As the hardware is component based, new hardware can be plugged into the existing system and servers can be configured quickly using scripts which prevents configuration drift by enforcing a common design. There is still complexity when moving between versions of hardware, but the scope is greatly reduced and once the migration path is determined, the roll out is greatly simplified.

The homogenization also aids with the patching strategy. As there are fewer variations in the systems, the same patch is applied more broadly, reducing the scope of the test and verification efforts and, as the patches are rolled through multiple environments, shortens the time it takes to get the patch to production. Figure 3 clearly shows the difference between a bespoke architecture and a homogenized one. The complexity and maintenance of the installation is significantly reduced.

The Sys Admins also realize a benefit to a homogenized environment. The knowledge acquired in one environment is directly applicable to other environments as they too are more homogenous and by reducing their workload, the Sys Admins can spend more of their valuable time performing more "value-added" tasks such as performance optimization and preparing for future updates.

NON-STANDARD COMPONENTS

Not all the systems used by NIF can be moved to homogenized components. Typically, this is driven by devices that are attached to computers that require a specific operating system and drivers version that are not available on more modern hardware or operating systems. In these cases, the cost of moving to new hardware may include the cost of buying a new instrument which makes the return on investment of the update prohibitive.

In these cases, the IT team needs to work with the stakeholder to explore alternative approaches.

Firstly, check that the customer has tried to use a new configuration or verify that they are not resistant simply because it will introduce change to their setup. In either case, working with the stakeholder with the update process can ease the migration creating a win for both parties.

If that is not possible, attempting to run the component inside a virtual machine can mitigate risks although this can be problematic if the component requires direct control of a device.

Finally, if there is no other option, isolate the host from the network so that the risk of running a nonstandard component is minimized as much as possible. It is also important to begin a review to determine the long-term strategy for managing the component and to communicate to stakeholders the risks associated with running in this way. If the component is of high value, ensuring that it is on a regular backup schedule is also recommended.

The process is a graded approach to handling nonstandard or unsupported components in order to minimize risk to the infrastructure and maximize the return on investment to the stakeholders

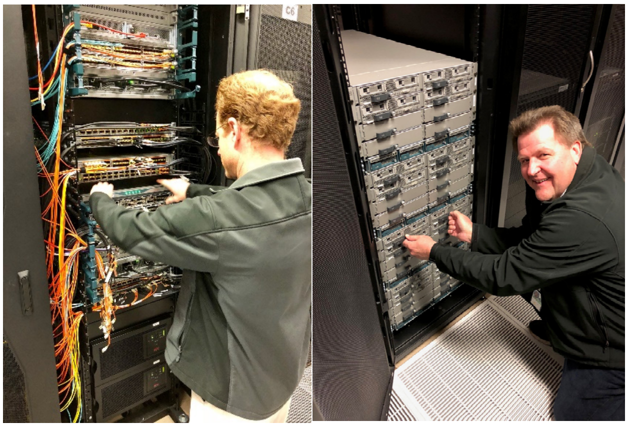

Figure 3: Before and after photos showing the difference homogenization makes to the infrastructure installation.

USING DATA TO MAXIMIZE ROI

Cost is a primary driver of when making decisions about modifications to the infrastructure. Understanding the use case of the update is essential to the process. There are many tools that can assist with this. Network, database, storage and compute all have tools that monitor the use, load and performance patterns of components and developing a thorough understanding of these patterns leads to more informed decisions that in turn maximize the return on investment.

When considering updates to storage such as adding more capacity or migrating data to the cloud, this kind of analysis is very applicable.

Users always want more storage and the data is always considered “hot” i.e. the user wants access to it immediately. Simply adding storage can be expensive especially if it is flash based for example.

A recent analysis of one of the drive arrays at NIF showed the usage pattern shown in Table 1.

Table 1: Drive I/O Usage Pattern

Usage Amount	Percentage
Zero	53%
Minimal	31%
Normal	14%
High	2%
Extreme	0%

Despite the users desiring immediate access to the data, 84% of it is rarely looked at if at all. Clearly migrating this data to high performance and high cost flash would not be a good investment. These types of files would be good candidates for cheaper storage solutions and possibly even the cloud particularly the cold storage and even deep freeze paradigms. However, care should always be taken when considering the cloud. The cost of data retrieval can be significant, and the buyer should be aware of these costs before committing to such a strategy.

Extreme use was driven by database access and at less than 0.5% of the total use, this is an excellent candidate for flash-based storage as improvements in the database performance tend to improve those systems that utilize it as well.

USING THE CLOUD FOR APPS

As computing paradigms have changed over the years, return on investment (ROI) has been driven by different factors. The monolithic applications of the 60s, 70s and 80s realized their ROI through automation and by doing things more quickly and greater reliability.

The next iteration led to a client-server distributed architecture and the use of the cheaper distributed computing enabled the ROI.

The third iteration was microservices. It is in the cloud where the industry would have us believe the ROI can be realized. However, this is not the case for many applications.

Each iteration is shorter than the predecessor, and the ability for large control systems and their supporting tools to realize the benefits of migrating to the new technologies is increasingly difficult. Even some private sector companies have not found the cloud to be as a great a cost reducer as they had anticipated.

In assessing the cloud as a place for legacy applications, it became clear to the NIF&PS IT team, that there were several recommendations that had to be considered before making the transition.

- Stay away from refactoring old applications that are built using very old languages and databases. The rework required to move it to the cloud, could be greater than the cost of simply keep it running on-prem.
- Stay away from applications that were poorly designed in the first place as they take a greater amount of work to migrate and maintain.
- Stay away from applications that are tightly coupled to the data store unless the data store is going to be migrated as well.

Once again, the decision comes down to the ROI. If the case can be made, and the user case supports it, going to the cloud may be a great idea. To date, no legacy NIF applications have been moved to the cloud.

DEVELOPING A DEV OPS PHILOSOPHY

Updating the mid-tier is essential because it is common vector for attack from the outside. However, simply demanding that the SW teams that use the platform update their code to meet the requirements of the IT team does not usually work.

A better approach is to develop a Dev Ops mentality and integrate with the SW team. This helps IT understand their issues; schedule pressure, resource constraints etc.

Providing the resources to do the work the IT team needs to have within the SW team leads to a win for all parties. The mid-tier team learns more about the application domain and how they are supposed to function, and they also get opportunities to code and hence an opportunity to learn new skills. The SW and IT teams get an improved security posture as the updates are rolled out through the various environments.

CONCLUSION

Several recommendations can be made from the experiences and learning moments over the last couple of years.

The IT Team

The team needs to ensure that updates are made with an understanding of the desired risk posture and the required return on investment. This may require updating processes and the skills to the team members.

Tools

To ensure that desired results are achieved, tools are needed to monitor the components of the system, to ensure that resources are utilized efficiently and to enable observability into the system. However, monitoring and data gathering should be reviewed and actionable.

Automation

Wherever possible use tools to reduce the burden on the team. Automation will provide better repeatability, maintainability and hence availability. It also allows the team to spend their time on value added tasks for the enterprise while growing their skills sets.

Stakeholders

Everyone from the sponsors to the cyber team to operations and the software teams needs to understand and buy into the evolution of the infrastructure. It makes understanding the risks, the trade-offs, the ROI and the reasons for change easier to communicate and hence value.

Evolving the infrastructure is difficult but it is not impossible if all stakeholders are working towards a common vision it.

REFERENCES

[1] Center for Internet Security, https://cisecurity.org/controls/cis-controls-list/

[2] CISCO, https://www.cisco.com/c/en/us/products/servers-unified-computing/index.html