Philips IntelliSite **Open pathology platform**

Dr. Bas Hulsken, CTO Philips Digital Pathology Solutions

Since the development of modern anatomical pathology as a medical discipline in the 19 Century it has been uniquely bound to the location of the specimen, requiring direct (microscopic) inspection of the biopsy. Today Digital Pathology is breaking those bonds, allowing the practice of anatomical pathology anywhere, independent from the physical location of the biopsy and thus the patient. Remarkably enough, the coming of age of high speed and high quality digital microscope slide scanners coincides with the success of globe spanning digital sharing and collaboration networks: the collaborative encyclopedia WikiPedia, GitHub for collaborative software development, Office 365 for business office software. Digital Pathology is enabled by these technology revolutions to convert local pathology labs into globe-spanning virtual pathology networks. In this white paper I will share with you Philips' vision on the future of collaborative digital pathology, and explain why it needs an open platform approach to succeed.

Why an Open Platform?

Digital Pathology needs to bring mobility and flexibility to the pathologist, so that he or she can work from anywhere, and present, share and participate fully in multidisciplinary team meetings, tumor-boards, and conferences. All without any concessions to patient safety and privacy. A fatal -- though perhaps understandable -- mistake would be to address this by simply replacing the current anatomical pathology's practice of moving physical tissue blocks and slides from A to B by a practice of moving digital image files of microscopy slides from A to B. There are three reasons why enabling collaboration in digital pathology by sharing whole slide image files will not be productive:

1. Security, privacy and patient consent are impossible to guarantee when patient records, such as whole slide images, are exchanged as files via network shares or portable media.

November 2016

2. Mobile workers require data that can be accessed from anywhere. The mobile software paradigm does not, and will never, support local storage of large data sets such as medical images.
3. The sheer sizes of digital whole slide images prevents effectively storing, sharing and analyzing on anything but server grade hardware that does typically not allow file based interactions.

If exchanging image files is not the answer, then what is? Well, we'll only have to look at how the aforementioned digital collaboration networks addressed this issue: Wikipedia, GitHub and Office 365 have all broken with the old practice of exchanging information on the basis of (standardized) files. Instead they make the creation, sharing and consumption of information an integral part of their digital platform. You do not upload a finished document to WikiPedia, you create it on the platform; you do not download a complete encyclopedia either, you view the page of the article that you are interested in right now.

The advantages of this approach are evident: your information is never outdated, you directly see the contributions of your collaborators, and all modifications and transactions on documents are done centrally on the platform, ensuring integrity, security and provenance of the information. Compare this to a whole slide image stored on a thumb drive, which does not allow access from your mobile devices, allows untraceable copies by anyone with access to the drive, does not protect against accidental or intentional data corruption or deletes, and won't actualize updates or corrections to the patient or case information. Understandably the practice of using standalone software working with documents and images on your local desktop is giving way to digital productivity and collaboration platforms where all documents are stored in the platform, either on the premise of the company or institute, or in the (private) cloud, and can be accessed by any authorized user, from anywhere, from any device.

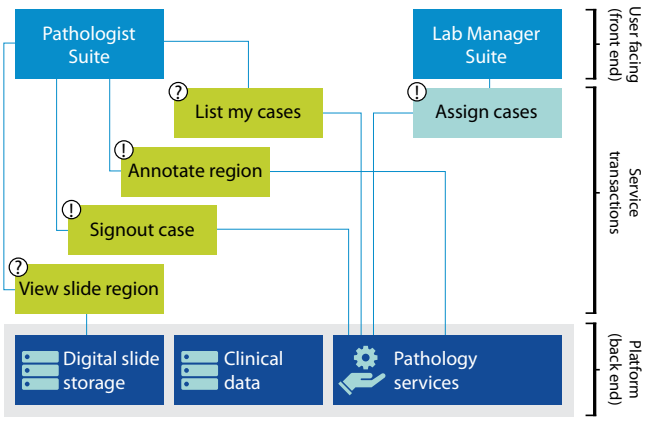

Figure 1: Simplified architecture diagram for an open pathology platform. Users interact only with the frontend applications, such as the Pathologist Suite to work cases. These frontend applications themselves are not responsible for storing or modifying medical data; instead they rely on databases and service brokers in the backend part of the platform. They are clients, communicating by simple requests with the backend server to retrieve or modify information in the databases or to start or stop tasks running in the service brokers.

Powered by Web 2.0 Technology

To bring the benefits of digital collaborative productivity platforms to pathologists, lab managers and researchers, an open platform is needed that leverages the technology that is powering today's online productivity platforms; regardless of whether the open pathology platform runs in the cloud, or on the hospital premise. The new internet build on these collaborative productivity platforms is often referred to as The Web 2.0[1]: an internet focused on usability, user generated content, and interoperability.

Figure 1 shows a breakdown of a simplified open pathology platform based on Web 2.0 technology. There are three essential elements:

1. HTML5 applications support working from anywhere, on any web browser equipped device. The user only interacts with these frontend applications.
2. Scalable and Highly Available backend services and databases reliably store data in the presence of unreliable frontend applications, and enforce security, adherence to patient consent, access privileges and data provenance.
3. RESTful communication interfaces[2] between the backend and frontend applications ensure robustness in the presence of multiple simultaneous users and unreliable connections and frontends.

The robustness and scalability of these platforms come from a large part from the uniform and simple RESTful style interfaces. For example an actual request to Philips IntelliSite Pathology Solution to view a slide image region as shown in Figure 1 has the following simple and human readable syntax:

```
GET   /ImageService/service.svc/image/slideid/pixeldata
      ?x=0&y=0&width=512&height=512&level=2
      HTTP/1.1
HOST  pathology.yourlab.com
```

Which is answered by the backend host (pathology.yourlab.com) with an equally simple, human readable, response containing a picture of the requested slide region of pixels at coordinate at 10X magnification (level 2):

```
                HTTP/1.1 200 OK
Content-Type    image/png
Content-Length  58643
Server          IntelliSite Open Pathology Platform/3.2.0
[PNG Image Data]
```

Due to the simple, human readable, communication interfaces of REST, Web 2.0 platforms are maximally transparent, and interoperability is easy. HTTPS connection encryption provides security and privacy. For the world wide web these simple interfaces avoided the slow process of standardizing complex binary document formats and enabled the rapid emergence of software ecosystems around open platforms, with contributions from industry, academia and the open source community. These open platform ecosystems are a big part of the success of the Web 2.0. But are by no means limited to the public web: Web 2.0 technology is equally suitable to power open platforms on local hospital IT systems.

The Philips IntelliSite Open Pathology Platform

The Philips IntelliSite Pathology Solution will be an open collaborative automated digital image creation, management and analysis platform that is modeled on the Web 2.0's open platforms. It's modular architecture is shown in Figure 2, with the components of the IntelliSite Pathology Solution in blue, the generic IT components on top of which the platform is deployed in yellow, and any third party components that would be deployed on the platform in green.

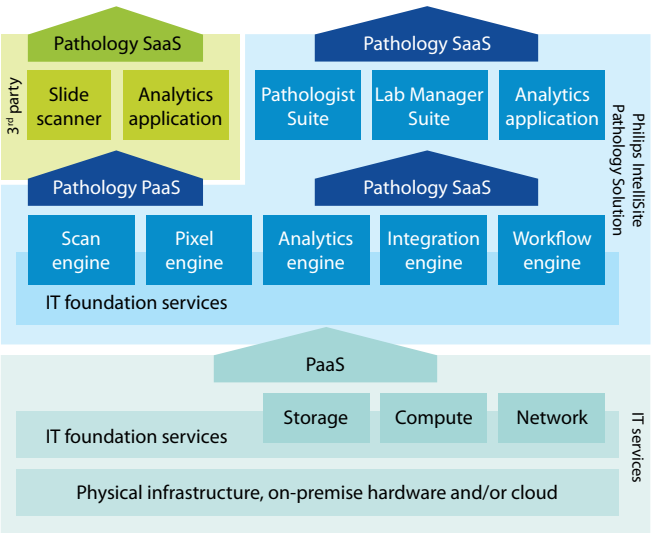

Figure 2: High level architecture of the Philips IntelliSite Pathology Solution (blue) as deployed on an IT platform (yellow). The communication interfaces between the different layers of the platform are indicated by the block arrows. The IT services are provided to the pathology platform by a Platform as a Service (PaaS) style interface. The pathology platform's services can be consumed by the applications using a Pathology PaaS interface. The Pathologist Suite and other frontend applications can be consumed as a Software as a Service (SaaS).

By designing the pathology platform to deploy on top of a generic IT platform, important requirements, such as scalability, availability and reliability -- which are requirements common to all large scale collaborative productivity platforms -- are delivered. An additional benefit is that by decoupling the pathology platform from the IT hardware, the pathology platform can be deployed on the hospital premise, in the cloud, or a mix of both.
As generic IT services are delivered by the underlying IT platform, the open pathology platform provides services that are specific to pathology. There are five main classes of services that a pathology platform should deliver in order to build meaningful solutions on that platform. In the Philips IntelliSite Pathology Solution, we have called these ``engines'':

1. The Scan Engine digitizes the slides; it is physically located in the scanner, but a functional part of the platform; apps will program scan protocols, and track it's status to ensure a seamless workflow.
2. The Pixel Engine provides all digital image services: display, processing, communication and storage. Digital pathology's large medical grade images require a specialized service, powered by the Philips iSyntax image format[3].
3. The Analytics Engine will be programmed by apps to execute analytics algorithms on whole slide images and associated data. The IT platform provides the computational infrastructure.
4. The Integration Engine provides standards-based interoperability with other enterprise information systems, such as the LIS, LIMS or EMR, leveraging Philips IntelliBridge broker.
5. The Workflow Engine is a programmable service that allows applications on the platform to define flexible workflows addressing the specific needs of the lab, the application or the hardware.

While these modular pathology service engines can be interacted with directly and individually, to guarantee security, safety and integrity, the platform can only be deployed as a whole; e.g. a stand-alone Scan Engine can not ensure correctness or privacy of the patient data without the presence of the Integration Engine.

Transparent Data Access

As an effective open pathology platform, the IntelliSite Pathology Solution caters to four basic needs of it's users:

1. Interoperability: allow 3rd party slide scanners and whole slide image viewers to seamlessly work with the platform, by providing simple, open but secure interfaces.
2. Flexibility: will allow migration of the data in the platform to or from alternative platforms using standards such as DICOM[4] and HL7[5].
3. Extensibility: will allow adding functionality to the platform by the customer or any other 3rd party to enable a pathology ecosystem.
4. Mobility: will allow working on, sharing and presenting content from any location, on any device.

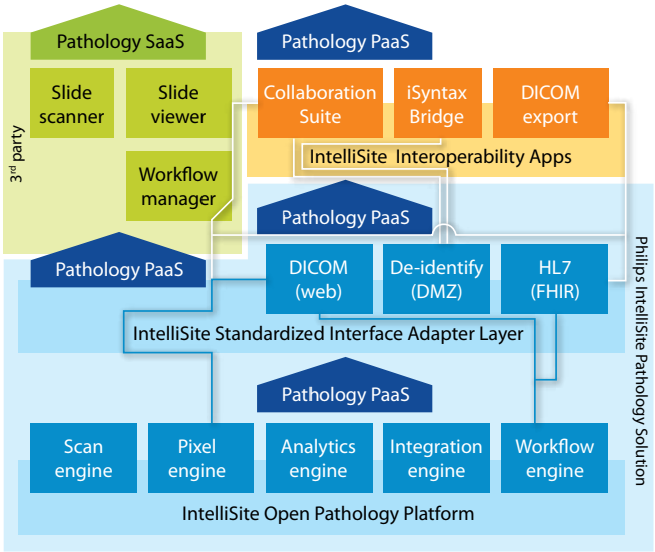

Figure 3: Philips IntelliSite Pathology Solution's open platform architecture providing scanner interoperability via the iSyntax Bridge, data migration via DICOM image export, and mobile sharing via the collaboration suite. The platform's functionality is exposed via standards compliant interfaces (DICOM, HL7), which can be used directly via the platform (PaaS), or indirectly via applications in the platform, such as the collaboration suite, iSyntax bridge and DICOM export apps (SaaS).

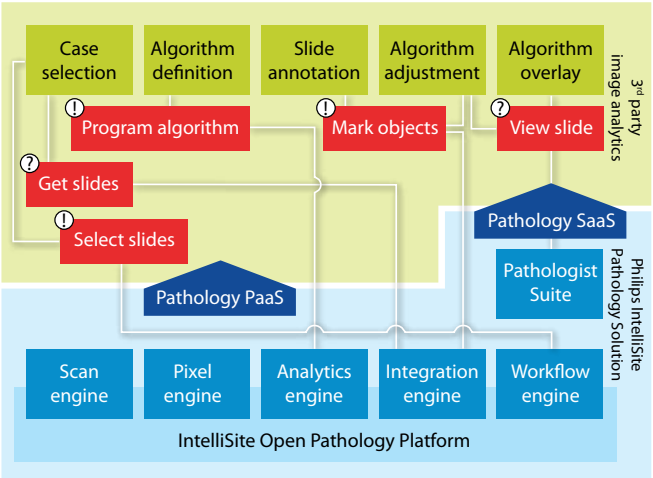

Figure 4: Example illustrating the deployment of a 3rd party app for image analytics (in green) on top of the Philips IntelliSite Open Pathology Platform. The RESTful, human readable, requests from the app to the platform are shown in red, where information retrieval requests are indicated with a question mark, and information storage requests are marked with an exclamation mark.

The Pathology Ecosystem

We stimulate 3rd party apps, from industry, academia or the opensource community to be developed and deployed on the IntelliSite Open Pathology Platform accessing its full and transparent data. Figure 4 shows a high level diagram of how e.g. an image analytics app would be deployed on top of the IntelliSite Open Pathology Platform. Using simple and human readable interfaces, the app would query the integration engine for the types of slides available, subsequently program the workflow engine to apply the actual algorithm -- programmed on the analytics engine -- to execute only on the relevant slides. An algorithm overlay method defined in the app would combine the algorithm results real time in the Pathology Suite to show the results to the pathologist. In this manner, 3rd party algorithms can be seamlessly integrated in the pathologist's workflow, and the computational burden of running the algorithm is on the shoulders of the underlying scalable IT platform as shown in Figure 2.

Philips HealthSuite

The Philips IntelliSite Open Pathology Platform will be part of the larger Philips HealthSuite, an open platform and ecosystem that brings together health data and medical history scattered over many apps, devices and systems in multiple places and formats, in a secure way. As a result, smarter and more meaningful connected health solutions for consumers, patients and care professionals will become a reality.

Bibliography

1. T. O'Reilly: \What is Web 2.0", O'Reilly Network (2005).
2. R.T. Fielding: Architectural Styles and the Design of Network-based Software Architectures, Ph.D. thesis (2000).
3. B. Hulsken: Philips' iSyntax for Digital Pathology, white paper (2016).
4. NEMA: Dicom: Digital Imaging and Communications in Medicine.
5. Health Level 7: HL7 Standards.

452220734711 - November, 2016